Spiking Neurons for Analysis of Patterns

High-performance pattern-analysis systems could be implemented as analog VLSI circuits.

NASA's Jet Propulsion Laboratory, Pasadena, California

Artificial neural networks comprising spiking neurons of a novel type have been conceived as improved pattern-analysis and pattern-recognition computational systems. These neurons are represented by a mathematical model denoted the state-variable model (SVM), which among other things, exploits a computational parallelism inherent in spiking-neuron geometry. Networks of SVM neurons offer advantages of speed and computational efficiency, relative to traditional artificial neural networks. The SVM also overcomes some of the limitations of prior spiking-neuron models. There are numerous potential pattern-recognition, tracking, and data-reduction (data preprocessing) applications for these SVM neural networks on Earth and in exploration of remote planets.

Spiking neurons imitate biological neurons more closely than do the neurons of traditional artificial neural networks. A spiking neuron includes a central cell body (soma) surrounded by a treelike interconnection network (dendrites). Spiking neurons are so named because they generate trains of output pulses (spikes) in response to inputs received from sensors or from other neurons. They gain their speed advantage over traditional neural networks by using the timing of individual spikes for

computation, whereas traditional artificial neurons use averages of activity levels over time. Moreover, spiking neurons use the delays inherent in dendritic processing in order to efficiently encode the information content of incoming signals. Because traditional artificial neurons fail to capture this encoding, they have less processing capability, and so it is necessary to use more gates when implementing traditional artificial neurons in electronic circuitry. Such higher-order functions as dynamic tasking are effected by use of pools (collections) of spiking neurons interconnected by spike-transmitting fibers.

The SVM includes adaptive thresholds and submodels of transport of ions (in imitation of such transport in biological neurons). These features enable the neurons to adapt their responses to high-rate inputs from sensors, and to adapt their firing thresholds to mitigate noise or effects of potential sensor failure. The mathematical derivation of the SVM starts from a prior model, known in the art as the point soma model, which captures all of the salient properties of neuronal response while keeping the computational cost low. The point-soma latency time is modified to be an exponentially decaying function of the strength of the applied potential.

Choosing computational efficiency over biological fidelity, the dendrites surrounding a neuron are represented by simplified compartmental submodels and there are no dendritic spines. Updates to the dendritic potential, calcium-ion concentrations and conductances, and potassium-ion conductances are done by use of equations similar to those of the point soma. Diffusion processes in dendrites are modeled by averaging among nearest-neighbor compartments. Inputs to each of the dendritic compartments come from sensors. Alternatively or in addition, when an affected neuron is part of a pool, inputs can come from other spiking neurons.

At present, SVM neural networks are implemented by computational simulation, using algorithms that encode the SVM and its submodels. However, it should be possible to implement these neural networks in hardware: The differential equations for the dendritic and cellular processes in the SVM model of spiking neurons map to equivalent circuits that can be implemented directly in analog very-large-scale integrated (VLSI) circuits.

This work was done by Terrance Huntsberger of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1). NPO-40945

Symmetric Phase-Only Filtering in Particle-Image Velocimetry

Performance is enhanced significantly with little increase in computation time.

John H. Glenn Research Center, Cleveland, Ohio

Symmetrical phase-only filtering (SPOF) can be exploited to obtain substantial improvements in the results of data processing in particle-image velocimetry (PIV). In comparison with traditional PIV data processing, SPOF PIV data processing yields narrower and larger amplitude correlation peaks, thereby providing more-accurate velocity estimates. The higher signal-to-noise ratios associated with the higher amplitude correlation peaks afford greater robust-

ness and reliability of processing. SPOF also affords superior performance in the presence of surface flare light and/or background light. SPOF algorithms can readily be incorporated into pre-existing algorithms used to process digitized image data in PIV, without significantly increasing processing times.

A summary of PIV and traditional PIV data processing is prerequisite to a meaningful description of SPOF PIV processing. In PIV, a pulsed laser is used to illuminate

a substantially planar region of a flowing fluid in which particles are entrained. An electronic camera records digital images of the particles at two instants of time. The components of velocity of the fluid in the illuminated plane can be obtained by determining the displacements of particles between the two illumination pulses.

The objective in PIV data processing is to compute the particle displacements from the digitital image data. In traditional PIV data processing, to which the

present innovation applies, the two images are divided into a grid of subregions and the displacements determined from cross-correlations between the corresponding sub-regions in the first and second images. The cross-correlation process begins with the calculation of the Fourier transforms (or fast Fourier transforms) of the subregion portions of the images. The Fourier transforms from the corresponding subregions are multiplied, and this product is inverse Fourier transformed, yielding the cross-correlation intensity distribution.

The average displacement of the particles across a subregion results in a displacement of the correlation peak from the center of the correlation plane. The velocity is then computed from the displacement of the correlation peak and the time between the recording of the two images. The process as described thus far is performed for all the subregions. The resulting set of velocities in grid cells amounts to a velocity vector map of the flow field recorded on the image plane.

In traditional PIV processing, surface flare light and bright background light give rise to a large, broad correlation peak, at the center of the correlation plane, that can overwhelm the true particle-displacement correlation peak. This has made it necessary to resort to tedious image-masking and background-subtraction procedures to recover the relatively small amplitude particle-displacement correlation peak.

SPOF is a variant of phase-only filtering (POF), which, in turn, is a variant of matched spatial filtering (MSF). In MSF, one projects a first image (denoted the input image) onto a second image (denoted the filter) as part of a computation to determine how much and what part of the filter is present in the input image. MSF is equivalent to cross-correlation. In POF, the frequency-domain content of the MSF filter is modified to produce a unit-amplitude (phase-only) object. POF is implemented by normalizing the Fourier transform of the filter by its magnitude. The advantage of POFs is that they yield correlation peaks that are sharper and have higher signal-to-noise ratios than those obtained through traditional MSF. In the SPOF, these benefits of POF can be extended to PIV data processing. The SPOF yields even better performance than the POF approach, which is uniquely applicable to PIV type image data.

In SPOF as now applied to PIV data processing, a subregion of the first image is treated as the input image and the corresponding subregion of the second image is treated as the filter. The Fourier transforms from both the first- and second-image subregions are normalized by the square roots of their respective magnitudes. This scheme yields optimal performance because the amounts of normalization applied to the spatial-frequency contents of the input and filter scenes are just enough to enhance their high-spatial-frequency contents while reducing their spurious low-spatial-frequency content. As a result, in SPOF PIV processing, particle-displacement correlation peaks can readily be detected above spurious background peaks, without need for masking or background subtraction.