Signal receiving section **130** includes a data receiving section **132** and a data storage section **134**. Data receiving section **130** may include a number of devices such as a modem and an analog-to-digital converter. A modem is a well-known device that comprises a modulator and a demodulator for sending and receiving binary data over a telephone line or other communication channel, while an analog-to-digital converter converts analog signals into a digital form. Hence, signal receiving section **130** may receive input signals "on-line" or in "real-time" and, if necessary, convert them to a digital form. As such, section **130** may receive signals from one or more devices such as a computer, a camera, a video recorder or various medical imaging devices.

The data storage section **134** serves to store input signals received by data receiving section **132**. Data storage section **134** contains one or more devices such as a disk drive, semiconductor memory or other storage media. These storage devices provide a method for applying a delay to the input signals or to simply store the input signals for subsequent processing.

In the preferred embodiment, the signal processing section **110** comprises a general purpose computer having a visual discrimination measure (VDM) **112**, a central processing unit (CPU) **114** and a memory **116** to facilitate image processing. The visual discrimination measure **112** can be a physical apparatus constructed from various filters which is coupled to the CPU through a communication channel. Alternatively, the visual discrimination measure **112** can be implemented as a software application recalled from the memory **116** and executed by the CPU of the signal processing section.

The signal processing section **110** is also coupled to a plurality of input and output devices **120** such as a keyboard, a mouse, a video monitor or storage devices including but not limited to a hard disk drive or a compact disk drive. The input devices serve to provide inputs (control signals and data) to the signal processing section for processing the input images, while the output devices serve to display or record the results.

The visual discrimination measure **112** assesses the visibility of differences between two sequences or streams of input images and generates an objective "just-noticeable difference" (JND) image metric. This metric can be expressed as a JND value, a JND map for each pair of input images or a probability prediction. In turn, the CPU may utilize the JND image metric to optimize various processes including, but not limited to, digital image compression, image quality measurement and target detection.

FIG. **2** depicts a simplified block diagram of the structure of the visual discrimination measure **112**, where two input image sequences **210** and **220** are processed to produce an image metric **250**. The visual discrimination measure comprises a temporal filtering section **230** and a spatial discriminator **240**.

In the preferred embodiment, the stimuli are two digitized sequences of images, input image sequence A **210** and input image sequence B **220**. For example, sequence A may comprise original images (e.g., a reference image sequence) while sequence B may comprise codec processed images of sequence A (e.g., a test image sequence). The input sequences represent time frames of sampled luminance distributions on a planar surface, i.e., as would be returned from a photometer sampling a uniform grid of closely spaced points on the surface of a display device. However, since the present invention is also designed to account for differences in the chrominance between the two input image sequences, the stimuli may include chrominance components as well.

Temporal filtering section **230** applies temporal filtering to both image sequences to produce eight separate responses (channels) **212–218** and **222–228**. For example, the luminance component (signal) of the input image sequence **210** are filtered into a low-pass temporal response **212** and a band-pass temporal response **214**. Similarly, the chrominance components (signals) of the input image sequence **210** is filtered into a low-pass temporal response **216** and a band-pass temporal response **218**. Thus, eight (8) channels are created for two input image sequences. The eight responses are received by the spatial discriminator **240** to produce an image metric **250**, which provides a measure of the visibility of differences between the two input image sequences. Furthermore, an optional retinal sampling section (shown in FIG. **3** below) can be incorporated to preprocess the input image sequences prior to temporal filtering, thereby further enhancing the predictive accuracy (relative to human performance) of the visual discrimination measure.

FIG. **3** depicts a detailed block diagram of the structure of the visual discrimination measure **112**. The visual discrimination measure comprises an optional retinal sampling section **330**, a plurality of temporal filters **335** and **334**, and a spatial discrimination section **240**. The spatial discrimination section (spatial discriminator) comprises a plurality of contrast pyramid filters **340** and **345**, a plurality of optional oriented filters **350**, a normalizer **355**, a gain control section **360**, a transducer section **370**, a pooling section (pooler) **375** and a distance section (distancer) **380**.

The retinal sampling section (retinal sampler) **330** receives and processes a plurality of stimuli to produce retinal sampled image sequences. Namely, the retinal sampling section resamples the stimuli to simulate the sampling process of a human eye. In addition, the retinal sampling section **330** may apply additional processings to prepare the image sequences for the effect of digital image processing. The retinal sampling section **330** is described in detail below with reference to FIG. **4**.

Referring to FIG. **4**, the retinal sampling section **330** comprises a frame rate adjuster **410**, a sampler **420**, a border inserter **430**, a smoother/interpolator **440**, a convolver **450**, a matcher **460** and a sampler **470**. The retinal sampling section **330** serves to approximate the sampling process of a human eye while adapting the images for optimal digital image processing.

The frame rate adjuster **410** receives the image sequence **405** (it may comprise input sequence **210** or **220**) and determines whether the frame rate is within the range that can be processed by a human eye. The frame rate adjuster incorporates a time-resampling method to address the possibility that the image sequence contains more frames per second than the vision model of a human eye can realistically process (e.g., 12,000 frames/second).

However, human vision can only resolve images with interframe interval greater than about 0.01 seconds, within which light intensity trades off against duration because the human eye is almost a perfect time-integrator. In cases of high frame rates, the frame rate adjuster **410** accumulates the averages (arithmetic means) of frames over consecutive rectangular time windows to arrive at new, subsampled pixel values. Namely, the frame rate adjuster **410** resamples an image sequence in time by evaluating mean pixel values for each pixel location over a number of frames N_t, and thereby