doi:10.1038/nature11150

1. System Comparisons

System	FoV (deg^2)	iFoV (deg^2)	Resolution elements (FoV/iFoV)	Etendue (m^2deg^2)	Optics volume (m^3)	System volume (m^3)
AWARE-2	6.0×10^3 [120° x 50°]	5×10^{-6} [2.3×10^{-3} ° sq.]	1.2×10^9	2×10^0	8×10^{-3}	3×10^{-1}
AWARE-40[1]	1.1×10^4 [120° circ.]	4×10^{-7} [6.2×10^{-4} ° sq.]	2.8×10^{10}	9×10^1	1×10^{-1}	5×10^{-1}
Multi-lens array[2]	4.1×10^3 [72° circ.]	1×10^{-5} [3.8×10^{-3} ° sq.]	4.1×10^8	2×10^1	1×10^{-1}	Cannot Estimate
Asymmagon[3]	4.0×10^3 [90° x 45°]	$< 6 \times 10^{-5}$ [$< 8.5 \times 10^{-3}$ ° circ.]	$> 6.7 \times 10^7$	3×10^{-1}	**8×10^{-3}**	**9×10^{-2}**
ARGUS[4]	2.8×10^3 [60° circ.]	2×10^{-6} [1.4×10^{-3} ° sq.]	1.4×10^9	**2×10^0**	**3×10^{-2}**	**3×10^{-1}**
LSST[5]	1.0×10^1 [3.5° circ.]	3×10^{-9} [5.6×10^{-5} ° sq.]	3.3×10^9	3×10^2	**3×10^2**	**1×10^4**
Pan-STARRS PS1[6]	7.0×10^0 [3° circ.]	7×10^{-9} [8.3×10^{-5} ° sq.]	1.0×10^9	2×10^1	**7×10^0**	**3×10^3**

Supplementary Table 1.1. Characteristics of gigapixel cameras. Quantities in bold were not reported (or were not directly derivable from quantities reported) in the references. These values are instead best estimates based on secondary information from non-quantitative information in the references (e.g. computer renderings). Systems in shaded rows are proposals only and have not been constructed.

Supplementary Table 1.1 provides a quantitative comparison between AWARE-2 and other gigapixel-class (or near-gigapixel-class) imaging systems. The majority of the values in the table are taken directly from public reports on the systems, or are calculated from reported values. Quantities in bold are estimates that rely on secondary information in the public reports (e.g. computer renderings which are assumed to be to scale). There are two items in the table that require special note: First, the Asymmagon is the sole comparison system that is film-based, rather than using electronic detection. This requires a different treatment of iFoV (and consequently, the number of resolution elements). In the electronic systems, the iFov is taken as the square, object-space angular extent of a single pixel. For the Asymmagon, the iFoV is modeled by a circular, object-space angular extent that is less than or equal to that which corresponds to the design MTF target of 30 cycles/mm. Second, the proposal for the Multi-lens array does not provide sufficient information about the detection, routing, and processing electronics to allow a reliable estimate of total system volume.

These systems segment naturally into terrestrial (AWARE and Asymmagon), airborne (multi-lens array and ARGUS) and astronomical (LSST and Pan-STARRS) classes. While the total number of resolution elements is roughly the same in the all classes, the astronomical systems have FoV and iFoV scales that are several orders of magnitude smaller than the terrestrial and airborne systems, while having volume scales that are several orders of magnitude larger.

LSST and Pan-STARRS are large infrastructure systems costing >$100M. Multiscale design yields no advantage for such large scale systems, where moderate field curvature can be compensated in precise sensor tiling and focal parameters are uniform and fixed. Multiscale design may, however, make wide field astronomical instruments available to those that cannot afford precisely integrated arrays. Multiscale systems provide modest microcamera cost, relaxed tiling tolerances and compensation for high radius of curvature focal surfaces. In the near term, however, terrestrial and airborne systems offer the most significant opportunities for multiscale systems. AWARE-2 matches ARGUS' pixel count in a 4x smaller optical volume. AWARE-40 exceeds the ARGUS pixel count by 20x with a 5x increase in volume. Beyond these comparisons, we note that to our knowledge the monocentric multiscale approach is the only lens design with the capacity to scale to 10-100 gigapixel systems. We also note again that parallel microcameras include parallel focus, exposure and image processing. This low-level integration of analog and digital image processing is a significant break in the long history, culminating in the Asymmagon's hybrid of analog acquisition and digital scanning, of analog photography and the concept of a continuous analog focal plane.

2. System Design

As illustrated in Supplementary Figure 2.1, AWARE-2 is a gigapixel multiscale camera consisting of a single objective lens and a multitude of microcamera lenses, each analogous to the eyepiece of a telescope. The objective lens is a refractive monocentric design with all spherical surfaces sharing a common center of curvature. Because the surfaces are symmetric with field angle, monocentric lenses have no aberrations except for spherical aberration, longitudinal chromatic aberration, and combinations thereof (e.g. spherochromatism). The design reduces these aberrations sufficiently so that the microcameras may correct residual aberrations. The focal surface is spherical with a radius of curvature approximately equal to the lens focal length, which for the AWARE-2 objective is 70.256 mm. The monocentric lens in the AWARE-2 design contains a nearly complete sphere of fused silica (n_d=1.458,V_d=67.82) surrounding the center of curvature of the objective. Surrounding this center sphere are two monocentric meniscus elements made of Ohara (Branchburg, NJ) S-NBH8 glass (n_d=1.720,V_d=34.71) of nearly equal outer radii of curvature. Unlike most monocentric lenses, there is no stop placed at the center of curvature of the lens.

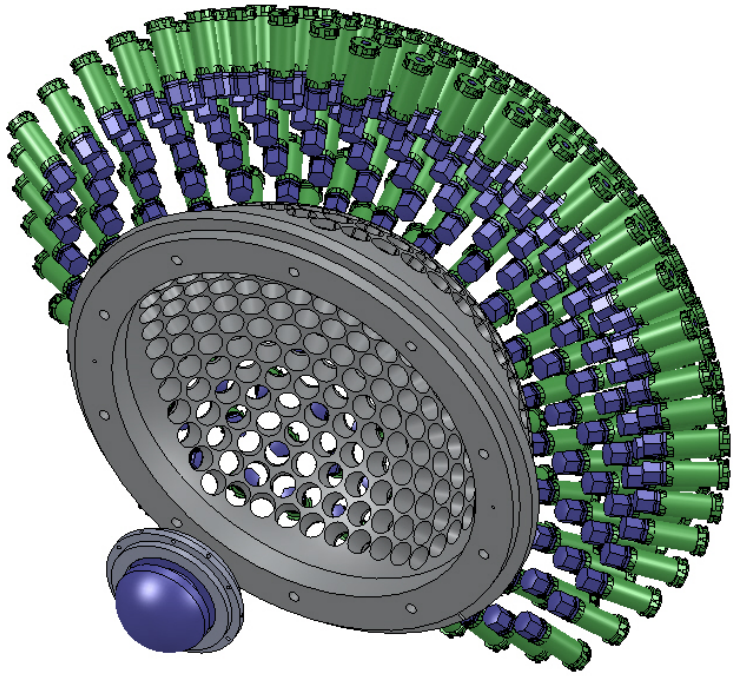

Supplementary Figure 2.1. An exploded view of the AWARE-2 camera. This shows the objective lens (lower left), the dome (gray in the center), and the dome fully populated with microcameras (green barrels and purple lenses).

An aluminum dome surrounds the objective and the intermediate image formed by the objective. The AWARE-2 dome was machined from a solid aluminum block using a 5-axis CNC machining tool and several annealing steps, including ports to fit up to 226 microcameras that cover a minimum of 120 degrees field-of-view for the entire camera. Each port is aligned such that the axis of the microcamera optics points directly towards the center of the objective. The objective is held in the dome by a three-armed flexure mount (not shown) so that it can be precisely positioned in space at the center of the dome. The arrangement of the cameras is a modified frequency-9, nonchiral, icosahedral geodesic[7] with the cameras axes aligned with the vertices, similar to the arrangement of carbon atoms on a fullerene molecule. Because a regular hexagonal packing of the sphere is not possible, a geodesic pattern is used to minimize the irregularity of the packing and maximize the overlap between microcamera fields. The largest spacing between microcameras is approximately 17% greater than the smallest spacing, achieving a high degree of packing uniformity. The center camera of the array is pentagonally-coordinated rather than hexagonally-coordinated and corresponds to a vertex of the original parent icosahedron used to construct the geodesic pattern.

The AWARE-2 microcamera design is shown in Supplementary Figure 2.2. The microcamera has an effective focal length of 11.68 mm, an f/# of 2.1, and a magnification of 0.487, and is placed so that the front surface vertex is 33.226 mm from the image formed by the objective. The field-of-view covered by each microcamera is 9.2 degrees of the overall camera field. The barrels are 37.25 mm long and 14.88 mm in diameter, and are machined from aluminum and black anodized to reduce stray light. The optics were inserted into seats machined into the barrel and secured with adhesive. The optical design of the microcamera uses four elements with all even-polynomial aspheric surfaces, made from molded polycarbonate (n_d=1.586,V_d=29.91) and Zeonex (Louisville, KY) E48R (n_d=1.531,V_d=56.04). Two of the elements form a cemented achromatic doublet lens of hexagonal cross-section, which is the large lens at the left of Supplementary Figure 2.2. The front polycarbonate element is of negative power, and the E48R element is of positive power. The hexagonal cross-section minimizes the vignetting that occurs at this surface. A microcamera was modified for the pentagonally coordinated camera at the center of the field of view. The hexagonal cross-section doublet can be thought of as analogous to the "eyepiece" of a conventional telescope, which approximately collimates the intermediate image. The remaining two elements and the sensor serve as the "eye" of the observer. The third element, the element second from the right in Supplementary Figure 2.2, is a positive-powered E48R element fabricated with a diffractive surface facing the hexagonal doublet. This element focuses the light onto the sensor, and the diffractive surface on this element corrects the secondary color because the two plastics alone do not have the diversity of dispersion needed to correct the color. Finally, the fourth small element at the right of Supplementary Figure 2.2 is a polycarbonate negative-powered element placed very close to the sensor. This element flattens the field and provides primary color correction. The stop is placed immediately before the third element at the diffractive surface. Because the stop is in the microcamera rather than in the objective, the entrance pupil of each microcamera is effectively rotated with the field angle of each microcamera, minimizing the vignetting inside the objective, as well as aiding in aberration control in the microcamera itself.

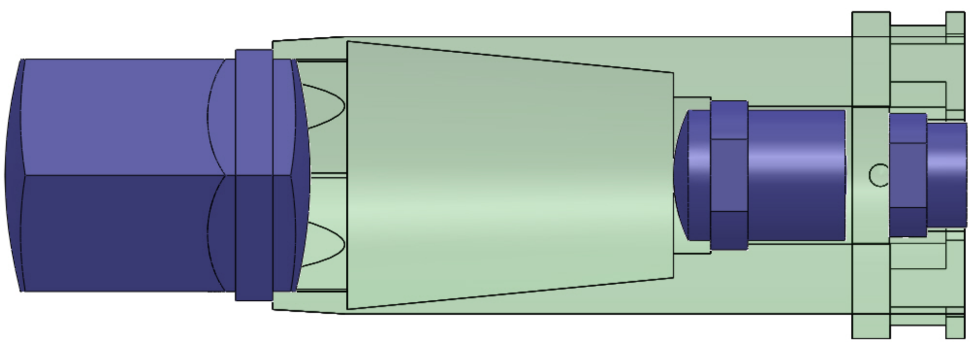

Supplementary Figure 2.2. CAD drawing of the microcamera barrel. The barrel is rendered semitransparent so the physical shapes and the mounting of the optics are clearly visible.

The sensor used was an Aptina (San Jose, CA) MT9F002 14 megapixel monochrome CMOS sensor. These sensors have a 4384 × 3288 array of 1.4 µm pixels, capable of 13 frames per second readout through its HiSPi interface. The 25° chief-ray-angle version of the sensor was used and the optics were designed to match. The image overfills the sensor slightly vertically (but not horizontally) to allow for efficient use of the sensor while still allowing for overlap between adjacent microcamera fields. Focusing of the microcamera is achieved by translation of the sensor using a lead screw and a servo motor. Translation of the sensor by 0.070 mm is sufficient to refocus the microcamera from infinity focus to 15 m range. The camera has an overall 16 mm entrance pupil size, corresponding to a 38 μrad resolution diffraction-limited Rayleigh criterion.

Radius	Thickness	Glass	Diameter	Comments
31.800	13.613	S-NBH8	59.394	Objective Start
18.187	18.187	F_SILICA	34.788	
Infinity	18.187	F_SILICA	33.904	
-18.187	13.613	S-NBH8	33.904	
-31.800	36.878		60.272	Objective End
	52.20087			Objective Focal Surface
42.1209	6.499998	S-BAH10	11.200	Microcamera Start
-42.1209	5.747543		10.94517	
11.79352	6.494901	S-LAH51	8.92721	
-14.9927	6.373718	S-TIH53	8.92721	
6.227476	1.8623		3.107624	
Infinity	0.7324965		2.501157	Stop
-5.095784	5.361673	S-FSL5	2.754262	
-4.204584	1.247128	S-NBH55	5.089383	
-6.628245	0.4993995		6.027518	
20.33398	2.007041	S-BSM16	6.800	
-20.33398	1.370708		6.800	
8.762839	6.504532	S-TIM5	6.856473	
8.348478	2.438452		4.884413	Microcamera End
Infinity	0.4	S-BSL7	4.5846	Sensor Window
Infinity	0.1		4.546649	
Infinity			4.532188	Image Surface

Supplementary Table 2.1. Prescription for a glass microcamera. The glass types are OHARA glasses and F SILICA is fused silica.

To benchmark performance of the plastic microlens, we also designed and built a second generation microcamera made from spherical polished glass elements selected from the OHARA catalog. The image is f/3 and telecentric and spanned 11.2 degrees field-of-view, as opposed to the 9.6 degrees field-of-view of the plastic microcamera. The prescription for the objective and microcamera together including the radii of curvature, element thicknesses and spacings, and diameters are shown in Supplementary Table 2.1. The design is based on a modified double-Gauss lens with modifications for image-side telecentricity and finite conjugates. The differences were included so that several different improvements could be explored in the next generation of microcamera, as well as examining the differences between plastic and glass optics. The glass microcamera was designed as a "drop in replacement" for the plastic camera including attachment of the sensor module. The glass microcamera is about 4 mm longer than its plastic counterpart. Unlike the plastic microcamera, which is designed to achieve approximately 40% nominal MTF at the 357 lp/mm Nyquist frequency of the sensor, the glass optics are designed for about 40% nominal MTF at 225 lp/mm. This is because the spherical elements did not allow for the control of the aberrations required to achieve higher performance while maintaining reasonable manufacturable tolerances when the number of elements were increased to have comparable nominal performance with the plastic microcamera.

A comparison of the spot diagrams of the plastic (left) and glass (right) microcameras are shown in Supplementary Figure 2.3. In each diagram, the dimensions of the plotted spot diagram are 0.010 mm on a side, and blue, green, and red symbols are the positions in the image plane for traced rays of 486, 588, and 656 nm wavelengths respectively. Four field positions are shown within each microcamera: 0°, 2.4°, 3.2°, and 4.6°. Three positions of the microcamera relative to the overall camera axis are shown, representing 0°-45° (over which there is little variation), 50°, and 55°. Vignetting inside the objective lens is responsible for the differences in performance at microcamera positions. A black ellipse in each diagram outlines the size of the diffraction-limited ellipse. In this case, the diffraction-limited spot is not radially symmetric due to tangential vignetting caused by splitting of the fields between adjacent microcameras. The performance of both of the cameras is nominally well below 0.1 waves RMS, so that the number of traced rays outside the diffraction-limited ellipse is relatively small, showing that the nominal performance is nearly diffraction-limited. Ray fans of these designs show that geometric rather than chromatic aberrations tend to dominate, especially at outer field angles, and that the field curvature and astigmatism are the limiting factors in image quality. The vignetting at the outer field angles servers to mitigate this problem so that diffraction-limited performance (albeit with a larger focal spot) is maintained throughout the field.

A given microcamera is constrained from imaging the fields of its neighbors by the angle subtended by the angle between its microcamera axis and those of its neighbors with respect to the objective center. Ultimately, this angle is limited by the physical size of the optics, the microcamera mounting hardware, and other possible interferences between microcameras. In this case, the optics bulk interference largely limited how closely microcameras could be packed, which had a minimum 7.0° spacing. To accommodate variations in spacing due to the geodesic configuration, the camera was designed to image up to 4.6° off axis. Angles beyond 3.5° off-axis imaged in a microcamera may be imaged closer to on-axis by another microcamera, so that the field beyond 3.5° is used to provide the overlap. Supplementary Figure 2.4 shows a ray trace of fields through two adjacent microcameras. Three points, A, B, and C, are imaged simultaneously in two microcameras at two different positions in their respective fields. In this case, field points more than 2.4° off-axis may be imaged on an adjacent microcamera, and fields closer to on-axis than 2.4° are imaged only by one microcamera. Because these fields are vignetted, reduced resolution occurs in the overlap regions, especially at the corner between three microcameras. Nevertheless, satisfactory performance is usually maintained out to 80% of the field edge.

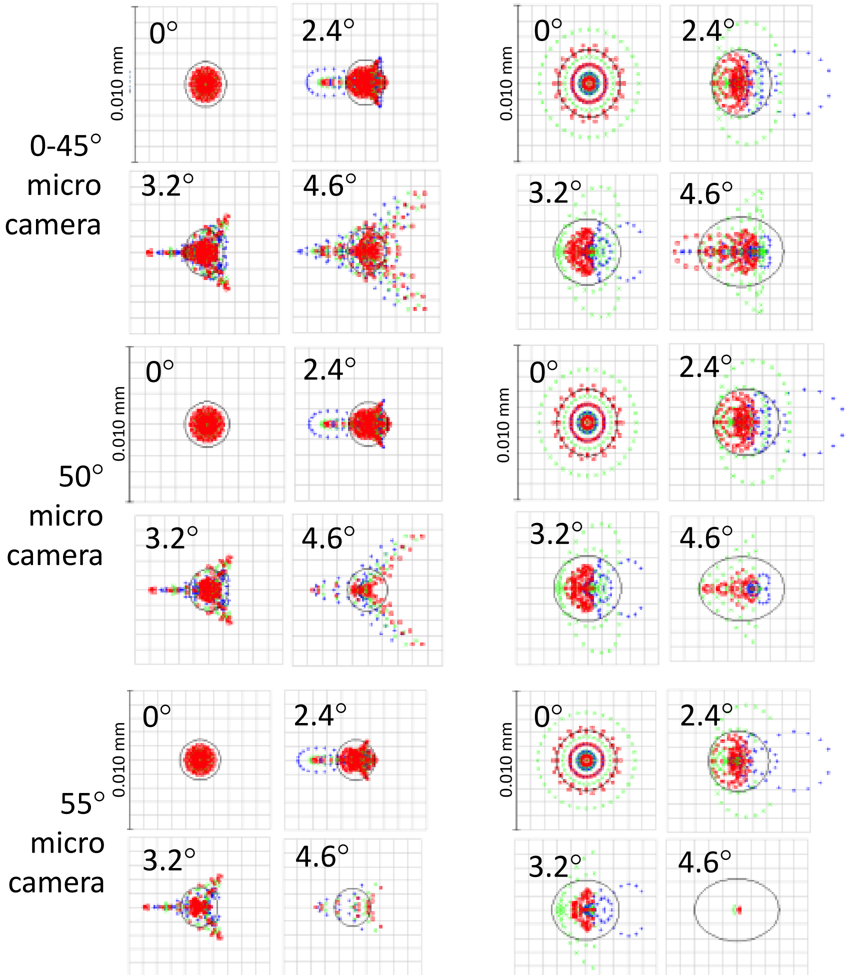

Supplementary Figure 2.3. Spot diagrams for plastic (left) and glass (right) microcamera designs. Each group of eight spot diagrams is labeled with the microcamera elevation angle with respect to the overall camera axis. Each spot diagram is labeled with its position inside the field of a respective microcamera, and has a dimension of 0.010 mm on a side. Red, green, and blue symbols denote traced rays for the wavelengths 486, 588, and 656 nm respectively. A circle is drawn on each figure denoting the diffraction-limited ellipse. The diffraction limited spot is elliptically shaped because of tangential vignetting. Also, the actual physical spot size is limited by diffraction when the ray traced spot size is comparable in size to the diffraction-limited ellipse size.

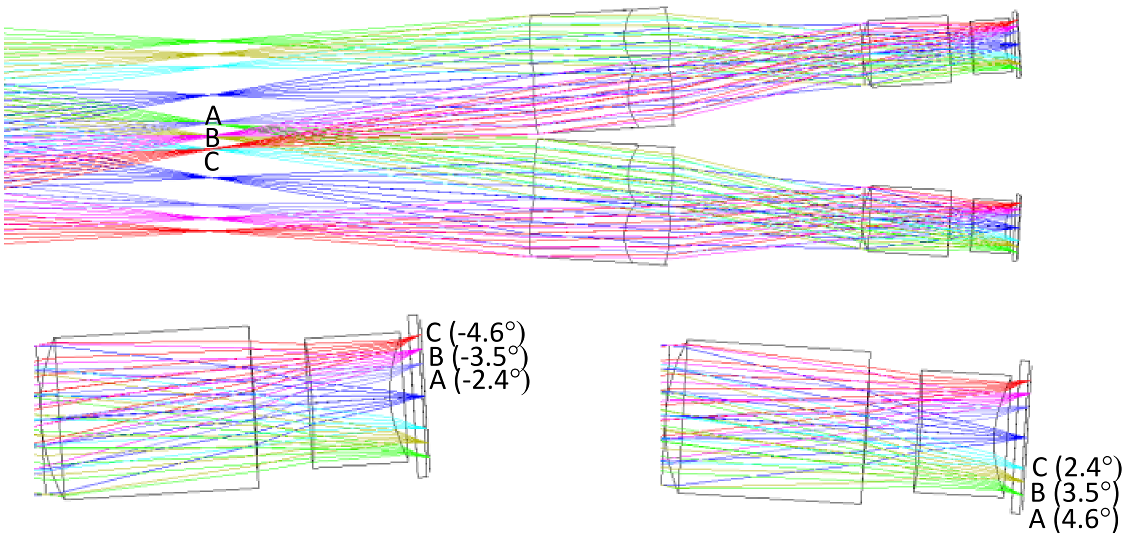

Supplementary Figure 2.4 A ray trace showing overlapping fields between microcameras. Three field points, labeled A, B, and C, are simultaneously imaged in two adjacent microcameras at two different positions in each microcamera. The points are labelled by the number of degrees off-axis these appear to be at in their respective fields.

The as-built AWARE-2 system is shown in Supplementary Figures 2.5-2.8. Supplementary Figure 2.5 shows the back of the microcamera mounting dome with optical barriers inserted. Clips hanging from the barrels are later used to retain the microcamera electronic package. Supplementary Figure 2.6 is a photo of the microcamera array looking into the mounting dome. Supplementary Figure 2.7 shows the gigagon objective mounted at the center of the dome. Supplementary Figure 2.8 is the AWARE-2 system with surrounding electronic control modules prior to mounting the gigagon.

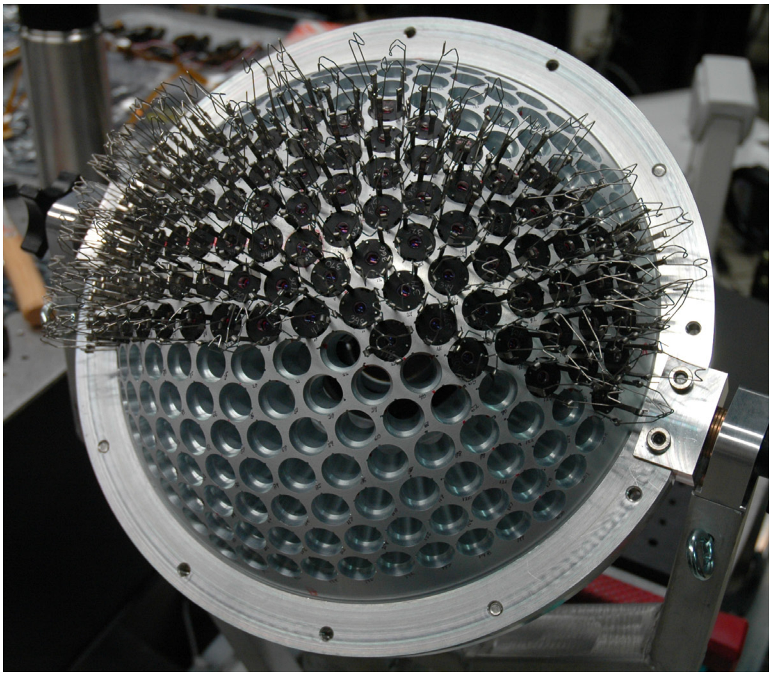

Supplementary Figure 2.5. Microcamera optics in the mounting dome for a 120x40 degree FOV AWARE-2 camera. The rear of the micro optic barrels are almost flush with the outside of the geodesic. The wire clips for attaching the microcameras are visible behind the barrels.

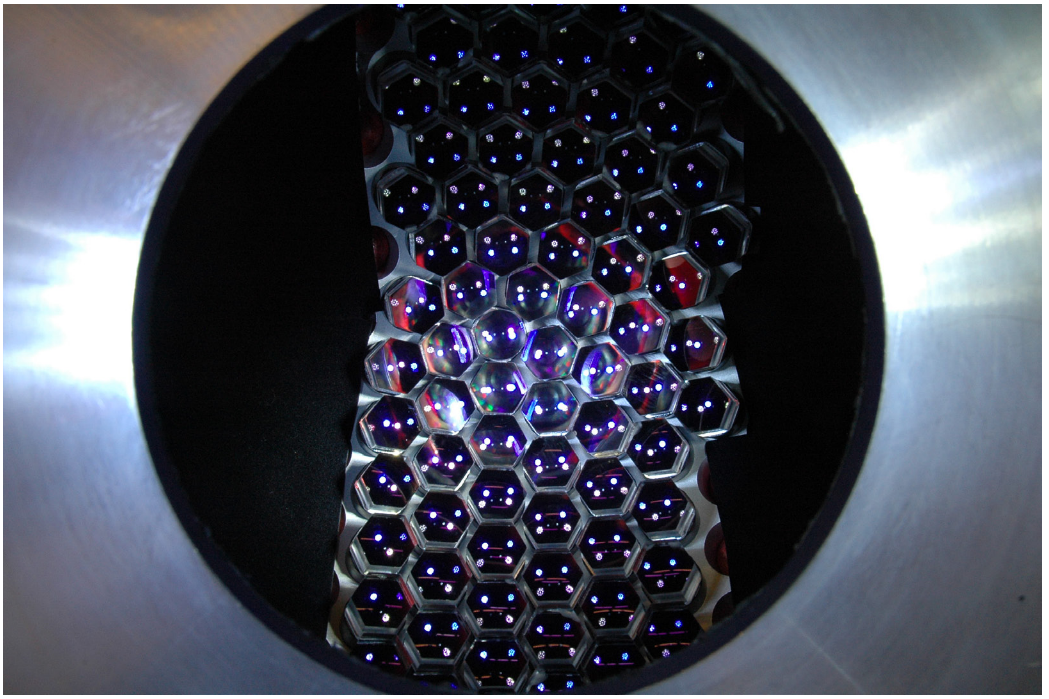

Supplementary Figure 2.6. Looking through the gigagon mount into the microcamera optics and blackened mount. The micro optics focus the intermediate focal plane of the gigagon onto the Aptina focal planes and correct aberrations from the gigagon.

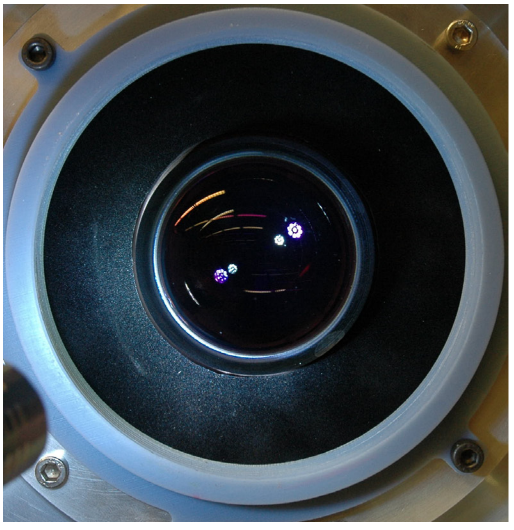

Supplementary Figure 2.7. The gigagon lens, mounted to the geodesic mount. The microcamera optics are behind this lens.

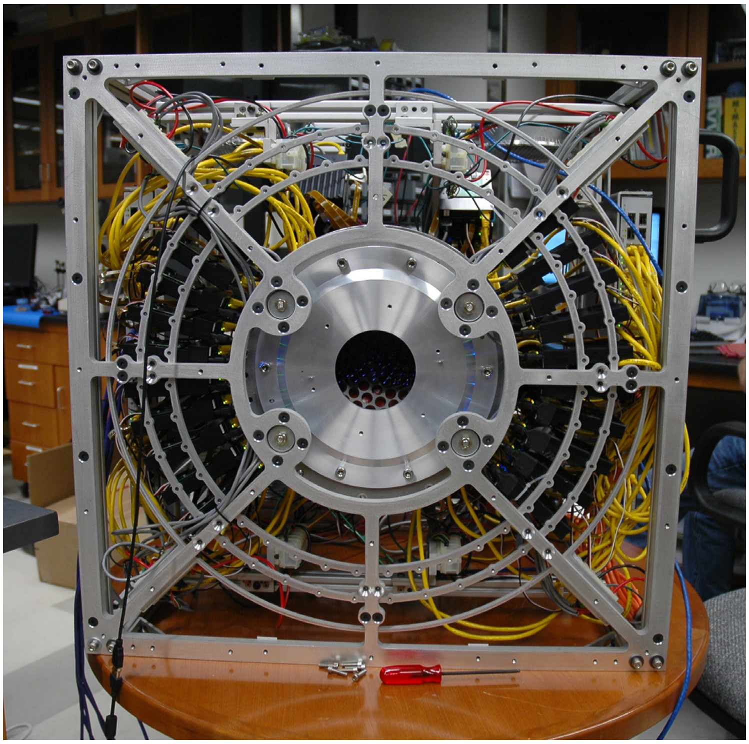

Supplementary Figure 2.8. The AWARE-2 system during assembly. The gigagon will be mounted in the hole in the center. Holes for the microcameras are visible through the gigagon mount. The electronic control modules can be seen mounted in the front, in a radial ring around the gigagon. Additional MCCMs are mounted to the rear. The frame is covered with PVC sheeting in the final assembly step.

3. Electronics and Processing

AWARE-2 supports up 220 microcameras each of which has Aptina's 14 megapixel monochrome (MT9F002) focal plane array (FPA) capable of providing 12-bit pixel data at 13.6 frames per second. This configuration presents a number of unique challenges for the electronics and image formation platforms. The electronics platform must manage data from many FPAs simultaneously and route this data to a processing system that can composite this raw multiscale data into seamless images for the user's display regardless of view position and scale. Further it was necessary to ensure that this architecture was flexible to allow for cameras with even larger pixel counts in the future. The system block diagram is shown in Supplementary Figure 3.1.

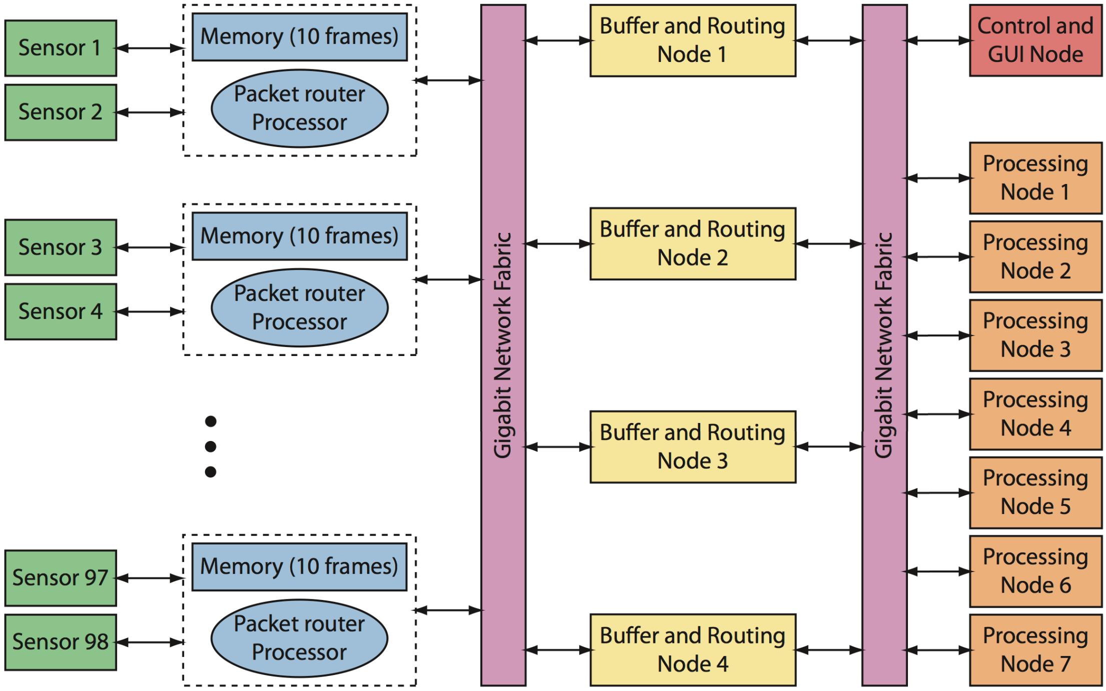

Supplementary Figure 3.1. Basic layout for the data flow and image processing architecture. The full image data lives near the sensor (FPA) and only image data required to satisfy the current operating mode is transferred downstream to the processing nodes. We view the (leftmost) gigabit network fabric as the dividing line between the imager electronics and the processing hardware.

The electronics platform was designed to support large numbers of microcameras along with scalable data throughput from gigabits per second up to over hundreds of gigabits per second. The primary building block of the electronics platform is the multiscale camera control module or MCCM. Each MCCM contains a single Altera Cyclone IV FPGA (EP4CGX150CF23) responsible for control and data management.

Data is routed from two microcameras via independent multi-lane LVDS serial inputs (HiSPi) allowing two simultaneous full size images (12 bit, 14 megapixel) to be received at 13.6 frames per second. The image data is then stored into 512 megabytes of mobile DDR memory. A single gigabit Ethernet link is used to provide both an output data channel and a bi-directional control channel. Image data can be requested either in whole or in part and will be subsequently streamed out over the Ethernet link at speeds greater than 800 Mb/s. Each MCCM and attached microcamera pair require 12 VDC and consume 6.7 W when operating at full capacity which, for the currently assembled camera, results in about 366 W/GP for MCCM with microcamera.

Additional features of the MCCM include:

- PWM outputs to position servo motors for focus control
- SATA links for high speed module interconnection
- USB 2.0 for auxiliary control and image transfer
- HDMI output for real time image display
- Firmware updates over Ethernet
- MicroSD slot for potentially storing image and configurations data

The FPGA firmware supports independent image pipelines for each microcamera along with a soft-core NIOS2 CPU. Each image pipeline provides hardware generated image metrics, buffering of image data into RAM, region of interest and scaling support, and hardware accelerated Ethernet packet generation. We have developed our own custom Ethernet protocol based on both TCP and UDP protocols to optimally allow for configuration and data transfer. The NIOS2 CPU allows for a C-like development environment for tasks that do not require manipulating image data at full speed. Tasks like MCCM configuration and automatic image control algorithms (exposure, gain, and focus control) fit within this functionality.

The AWARE-2 camera allows for a total of 110 MCCM units that are connected together with an appropriate number of gigabit Ethernet switches. A suitable number of ports on each Ethernet switch are reserved to exchange control and data packets with the data processing system. Scalability is achieved by simply changing the number of connected Ethernet cables. In the extreme case, allowing for one Ethernet cable per MCCM bypasses the need for a switch entirely. For most applications far fewer links to the outside world are needed. Display technology typically supports a limited image size and a user is only able to look and make use of a limited amount of data. Capping the throughput somewhat arbitrarily to support ten users each with their own HD display results in a required system output throughput of 1.6 Gb/s, which is possible with just two-gigabit Ethernet connections, a very reasonable number. This illustrates that there is fundamental high-bandwidth and low-bandwidth division within our camera where we have crafted the current architecture of the system to keep the high-bandwidth section as close the FPA as possible. In doing so we ensure that only the pixels the user wishes to see will bare the full cost of transmission. This requires the MCCM to be quite flexible in how image data can be requested to allow for arbitrarily sampled and positioned data that will form the eventual composited image for the user's HD display. We call this our 'window-view' mode of operation. We have also allowed for a snapshot mode that will eventually operate in conjunction with the live mode where we can freeze a specific image buffer across all MCCM devices and have that image set transferred out as bandwidth allows.

There are several features under active development that will further increase performance while reducing size and power consumption. One example is the previously mentioned SATA links, which are essentially 3.125 Gb/s serializer-deserializer (SERDES) channels. These links are designed for interconnecting MCCM units together and would replace the Ethernet switches when functional. This will save nearly 1W per sensor if one considers the onboard Ethernet circuitry as well as the power for the Ethernet switch. Ethernet ports will then be provided only at required points in the system where data must be transferred off the camera. Additional improvements will come from a revision 2 design and increasing the number of microcamera channels per MCCM to around 6. Practical wiring constraints as well as FPGA cost and complexity place a limit somewhere between 4 and 8 channels per MCCM.

As with the electronic architecture, a major challenge to image formation and processing in the AWARE-2 system arises from the vast rate at which data can be generated. To address this challenge, the operating modes and image formation pipeline have been designed from the ground up for the specific environment of a multiscale imaging system like AWARE-2.

AWARE-2 image data management is implemented on a 12 node computing cluster. Each node contains 12 GB of memory and is equipped with a twelve-core Xeon CPU running at 2.67 GHz. One node runs the graphical user interface (GUI) as well as the scheduler for the parallel processing jobs. Four of the nodes provide data routing and buffering to manage the data flow between the camera electronics and the processing workers. The remaining seven nodes are dedicated to the image formation task. Of the 84 available cores in these nodes, 81 are identified as parallel workers and assigned portions of the image to process.

Two operating modes are implemented on the system—'full-frame snapshot' and 'window-view'. In the full-frame mode, the system produces a full-resolution, full-field-of-view image with an angular resolution comparable to the native pixel resolution in the underlying focal plane arrays. In the second operating mode, 'window-view', the operator is presented with an approximately one-megapixel 'window' into the system imagery. The user may pan/zoom this window at will as the imagery data continues to update. Functionally, the two modes provide archival and exploratory capabilities, respectively.

The AWARE-2 image formation pipeline is responsible for correcting and fusing the individual microcamera images into the overall system image (or the currently-required image window). To facilitate the eventual system frame rates, the pipeline has three important features. First, compositing is done 'blind'—for image fusion, the geometric and optical parameters of the system are assumed known and the output image is computed in a non-iterative fashion. Second, the compositing algorithm itself is cast in a highly parallelizable MapReduce framework[8]. Third, estimates of the geometric and optical parameters are periodically updated through a registration process that can be amortized across many frame times.

Processing times on the system are not deterministic as a result of variable latencies in the operating system and the communication layers. In general, we observe that a full-frame snapshot (~ one gigapixel) takes approximately 180 seconds to generate (~5.5 megapixels/s). We observe that the live-view window (~ one megapixel) takes approximately one second to generate (~1 megapixel/s). While the core processing algorithms are O(N) in time (and memory), fixed overhead results in slight nonlinearity in pixel throughput as a function of pixel number.

The compositing (fusion) algorithm is divided into two parts, corresponding to the 'map' and 'reduce' steps of the map/reduce framework. In the map portion, each pixel measurement represented by a tuple of the form

((microcamera number, pixel location), pixel measurement)

is converted into a tuple of the form

((Θ_x, Θ_y), (pixel measurement, relative illumination)),

where Θ_x and Θ_y are the horizontal and vertical angles of the object space direction with respect to the overall system optical axis.

The conversion from camera and pixel location to angular location in object space is done using a low-order parametric model of each microcamera. This model is currently built using optical simulations of the prescription in the ZEMAX optical modeling environment, although ultimately we will use per-camera models developed from precise calibration. By design, relative illumination rolls off sharply at the edge of each microcamera field. A central region deemed to have sufficient illumination containing approximately 12 megapixels is used in image compositing. Pixel data is supplemented by an estimate of the relative illumination at the pixel location (the microcameras have significant vignetting as a result of the optical constraints of a multiscale system). The relative illumination is computed from a low-order parametric illumination model. This model is currently generated from global flat-field calibration measurements. The map step for each tuple is independent, and therefore trivially parallelized.

In the reduce portion of the compositing, the tuples that correspond to the same (discretized) location in object space are combined via a maximum-likelihood (ML) computation that identifies the most likely source irradiance at that location, $i(\Theta_x, \Theta_y)$, given the measurements, their expected relative illumination, and the belief that the measurements are corrupted by additive, white Gaussian noise ,

$$\mathrm{ML}\left[i(\Theta_x, \Theta_y)\right] = \frac{\sum_{k=1}^{N} r_k\, m_k}{\sum_{k=1}^{N} r_k^2}.$$

Here N is the number of measurement pixels associated with the specific discretized location in object space, r_k is the relative illumination for the k-th measurement pixel, and m_k is the pixel value. The reduce step for each location in object space is independent and is therefore also trivially parallelized. The net result of the reduce step is a list of tuples of the form

$$\left((\Theta_x, \Theta_y), \mathrm{ML}\left[i(\Theta_x, \Theta_y)\right]\right),$$

which represents the desired estimate of the source (i.e. the final image). For images taken with a global exposure setting, the result is the final 8-bit monochrome image. For images taken with varying microcamera exposures (i.e. for snapshot HDR imagery), the result is a 32-bit monochrome image. As displays typically have 8 bits of gray depth, we are then faced with a choice of how to map the 32 bit data into the 8 bit display space. Our current implementation allows for linear and logarithmic mappings as well as our version of tonemapping (a common nonlinear mapping that uses the majority of display dynamic range on shadows and highlights—this compresses the midtones and results in a display image that appears more normal to the human visual system.

The compositing algorithm relies on geometric and optical knowledge about the system. The optical knowledge is embedded in the low-order parametric models, while the geometric knowledge is represented as a triplet of angles for each microcamera—pointing (two angles, indicating the orientation of the microcamera optical axis with respect to the system optical axis) and clocking (rotation of the camera about its optical axis with respect to some global origin). There are 'as designed' values for these parameters, but manufacturing tolerances, vibration, and thermal drift invariably result in some deviation from the expected values. The values of these parameters are periodically updated by the system through the use of an automated registration procedure. In registration, the microcamera images are processed using the scale-invariant feature transform (SIFT)[9] and sped up robust feature (SURF)[10] algorithms to generate a list of feature-descriptors. This list is winnowed to a set of feature descriptors that appear in multiple microcameras and a global correspondence error is calculated using the current geometric and optical description of the system. Pointing, clocking, and magnification are then adjusted using a nonlinear optimization routine to minimize the global error. The geometric and optical description of the system is then updated with these new values.

4. System Performance

A. Example Images

A map of a stitched AWARE-2 image with the camera numbers overlayed is shown in supplementary figure 4.1. Each camera in this image was independently focused and exposed, then stitched into a single composite.

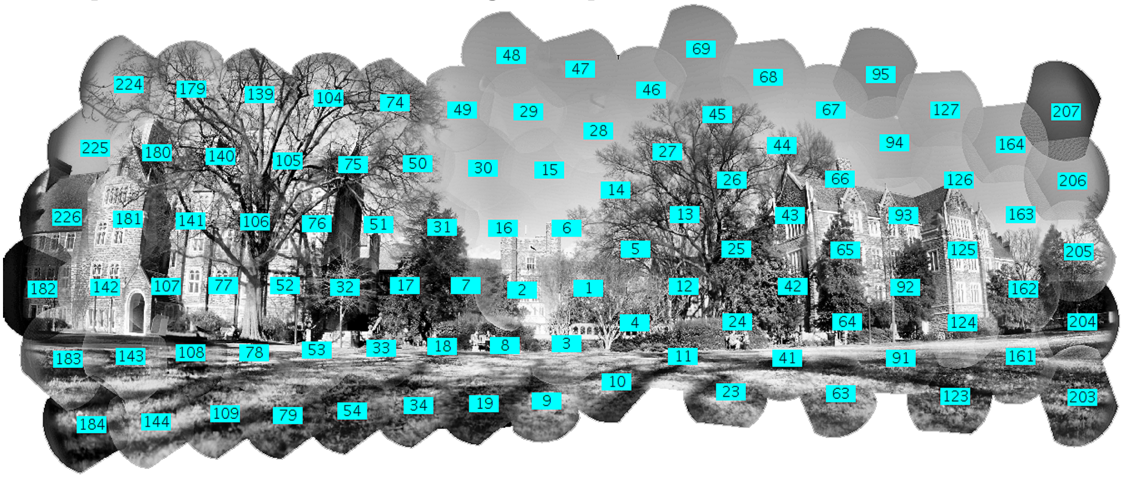

Supplementary Figure 4.1. Camera map for AWARE-2. This shows 98 microcameras with a 120 degree horizontal FOV and a 50 degree vertical FOV. This composite is an HDR, single frame stitch on the Duke University campus. The HDR image is synthesized from each camera having an optimal exposure, then being toned mapped into a single composite HDR image.

AWARE-2 requires a 12 volt, 37 ampere DC power for the MCCMs and cooling fans and a nominal 120V AC line to power the two Ethernet switches. An efficient AC-DC switching supply powers the DC circuits using a single 60A module, expandable up to three in a single 1U rackmount unit. The external supply to the camera is then a single AC power cord, capable of 120V AC at less than 5A. For conditions that require a smaller FOV, small banks of MCCMs can be powered on or off via software, reducing data transfer time and power consumption.

Low resolution (screen level) stitching can be done on a single core processor in a few seconds for on-site image analysis. Each image requires about 1.5 gigabytes of storage, making a 1 terabyte hard disk sufficient for a full day of fieldwork. Data download from the 49 MCCMs requires about 18 seconds.

Field images were captured at the Pungo Lake viewing platform in the Pocosin Lakes National Wildlife Refuge in North Carolina, USA on 5 December 2011. The camera was placed on a wildlife lookout 12 feet above ground level, where minimal interference from surrounding foliage obstructed the Lake. Although the AWARE-2 camera is capable of tremendous data output, reduced data loads and power requirements may be implemented in field work. A standard laptop computer is capable of operating the system, especially with the use of a USB powered mobile disk drive. A single gigabit Ethernet line was connected to the camera for data transfer, limiting frame rates, at the advantage of only needing a single node for capture.

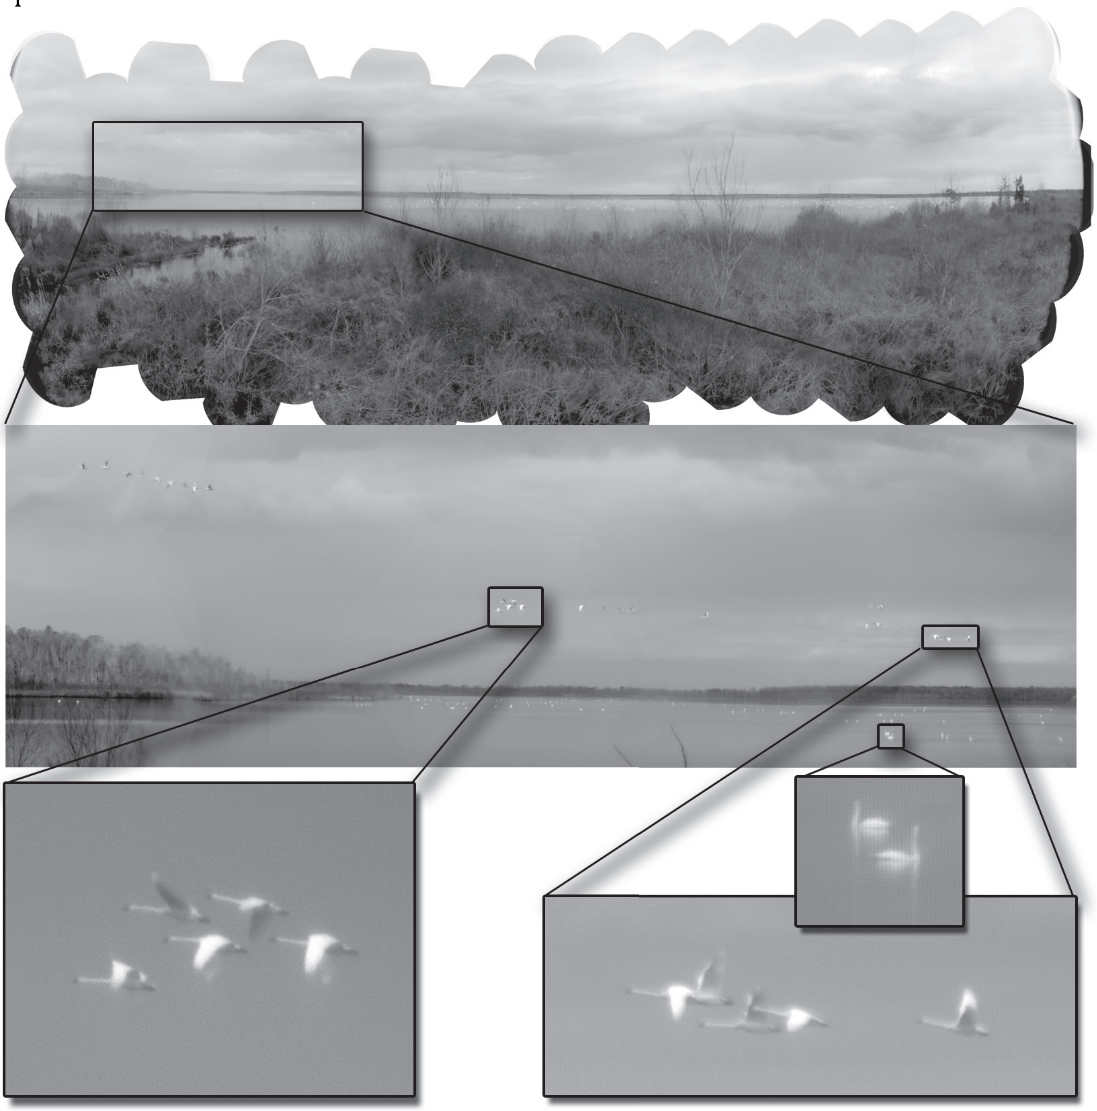

Supplementary Figure 4.2. Relative zoom levels on an image from the Pungo Lake viewing platform. The top image is the entire 120x50° FOV. The middle image is magnified by 4x and the bottom images are magnified by 32x.

The tundra swans on Pungo Lake ranged from 150 yards to over a mile away from the camera location. Uniform exposure and focus settings were applied on all 98 microcameras for the Pungo Lake images. Capture commands are sent simultaneously to all cameras but 50 ms uncertainty in the capture time leads to potential motion ambiguity between microcameras. Supplementary Figure 4.2 shows the entire FOV for the AWARE-2 camera, with a 4x zoomed in region, and three 32x zoomed regions of swans.

The AWARE-2 camera is well suited to taking pictures of dynamic, large FOV scenes. Supplementary Figure 4.3 is an image of conference attendees in front of the Carolina Theatre in Durham, North Carolina. This image shows both the dynamic range and synchronization of the microcameras across the FOV. Additionally, since each camera can focus independently, the full range from quite near to infinity is optimally focused, while still being an f/2.17 imager.

Supplementary Figure 4.3. The Carolina Theatre in Durham, North Carolina. Top shows a full, stitched image from the AWARE-2 camera, featuring 98 microcameras. The bottom insets show zoomed in regions of the image at 11.5x relative to the full image. Each camera has an optimal exposure, dramatically increasing the usefulness of the camera in high-contrast settings.

B. Resolution and MTF

Supplementary Figure 4.4 illustrates system resolution and degradation across the microcamera field. Supplementary Figures 4.4 (a)-(c) are images of 1951 USAF resolution charts located at 0°, 2.86°, and 3.91° with respect to the optical axis of microcamera number 17, which is centered at 21.1° with respect to the system optical axis. This represents 0%, 59.3%, and 81.1% of the full-field (defined as the central region of the microcamera image where the relative illumination is deemed sufficient). Supplementary Figures 4.4 (d)-(f) are similar images from microcamera number 142, which is centered 56.4° from the system optical axis. Camera numbers are indicated across the field in Supplementary Figure 4.1. Camera 142 is near the edge of the field, while 17 is near the center. Resolution charts were located at 0° (0% of full-field), 2.37° (49.2% of full-field) and 4.3° (89.2% of full-field) with respect to the microcamera optical axis. At the centers and approximate midpoints of the microcamera field, each image resolves group 2, element 3 of the resolution chart, which is at the Nyquist (1 bar feature per pixel) limit for the sensor. Resolution degrades to 1.4 pixels per feature in 3.1(c) and 2 pixels per feature in 3.1(f). Computational strategies to combine features from adjacent microcameras to improve resolution and relative illumination are under investigation.

Camera 17 was taken in afternoon sunlight from 69 yards, while camera 142 was taken under natural diffuse sunlight at 69 yards. The camera array was rotated three times to capture the 8.5x11 inch Air Force target in three different field locations for each microcamera. Supplementary Figure 4.5 shows a superposition of the three images for camera 017.

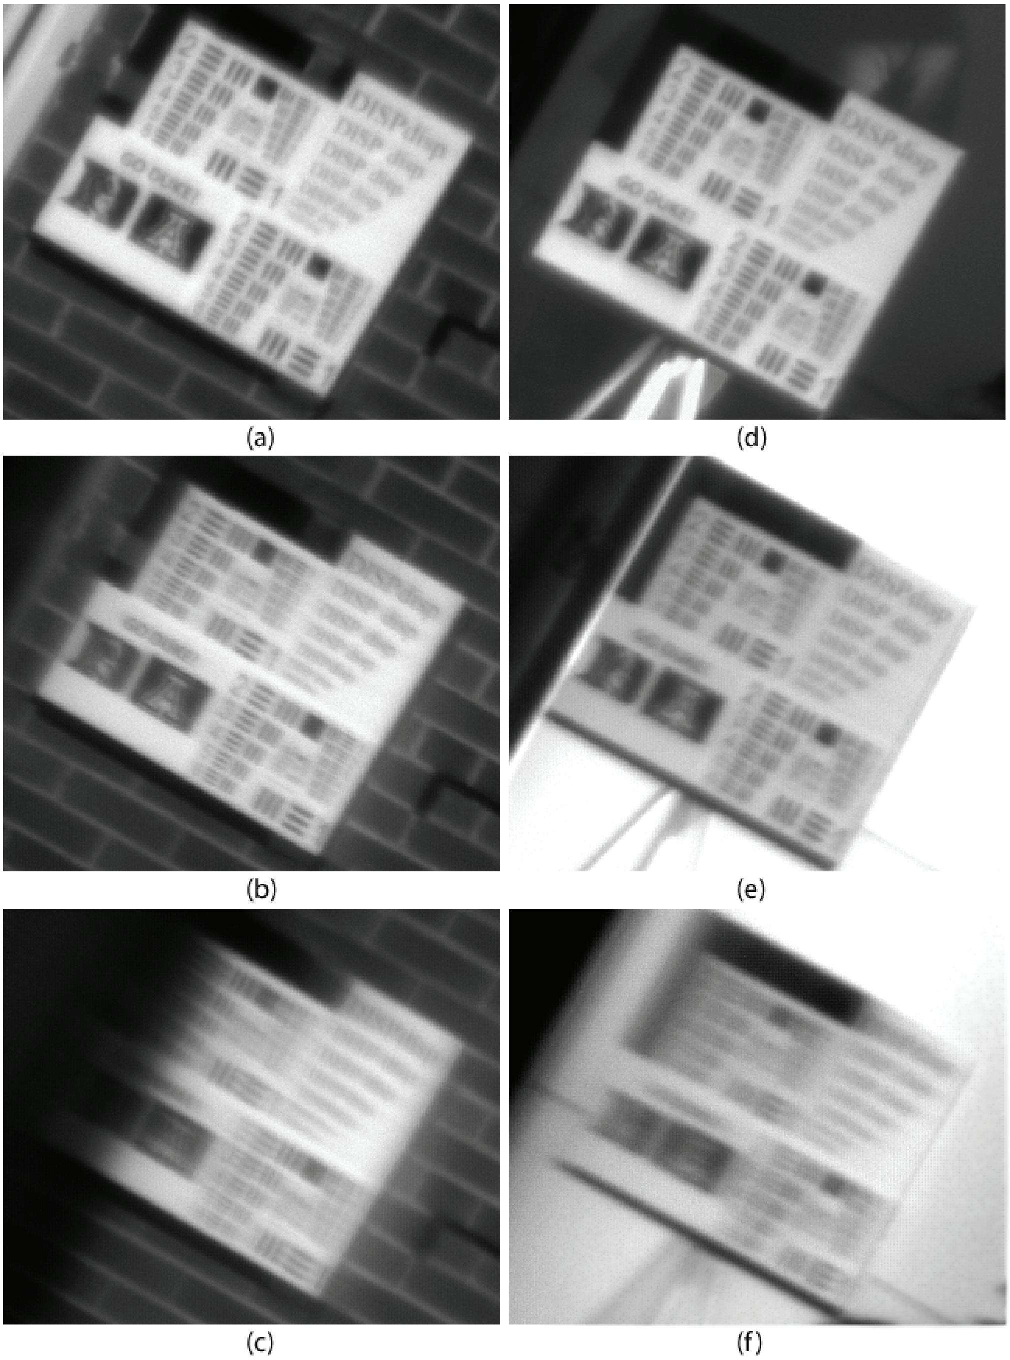

Supplementary Figure 4.4. Resolution test images at the center, middle and edge of the field for microcameras 17 (a-c) and 142 (d-f), using the original plastic microcameras. This figure shows the degradation in resolution from the center to edge of the field, caused in a large part by birefringence in the plastic optics.

Supplementary Figure 4.5. Superposition of three separate images, showing the field locations in camera 017 for the resolution charts shown in Supplementary Figure 4.1. The target was stationary on the building while the camera was rotated for the three field locations.

AWARE-2 system performance was also evaluated by measuring the modulation transfer function (MTF). MTF plots were generated by imaging a sharp black and white edge. Edge sub-images were processed using the commercial software Imatest. All images were unsaturated in both the dark and light regions. Imatest documentation was closely followed for particulars in acquisition and region of interest for the knife edge. Supplementary figure 4.6 shows the MTF in line pairs per millimeter (lp/mm) at the center of the sensor for a sampling of microcameras.

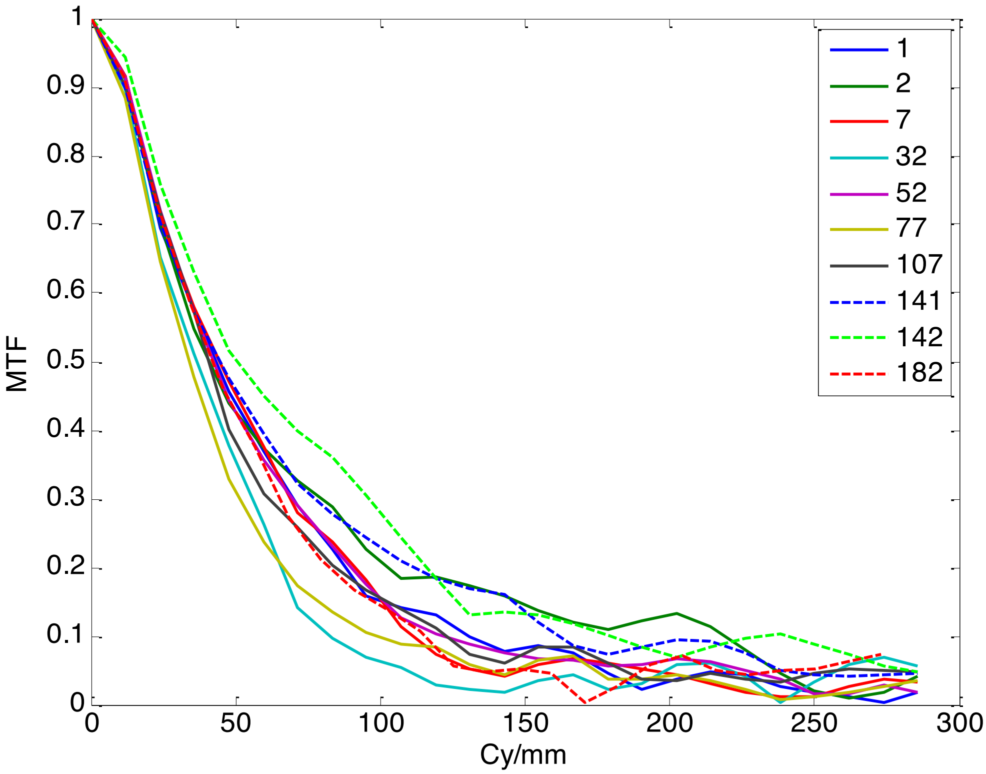

Supplementary Figure 4.6. MTF plots for 10 different microcameras throughout the field of the gigagon of the AWARE-2 camera. All plots are of the center field of each microcamera. Camera 1 is at the center of the gigagon, and camera 182 is at 60 degrees horizontal. Refer to Supplementary Figure 4.1 for details of the microcamera number mapping.

Microcamera number	1	2	7	32	52	77	107	141	142	182
MTF	16%	23%	18%	7%	18%	11%	17%	24%	30%	15%

Supplementary table 4.1. MTF at 95 lp/mm for Supplementary Figure 4.7. The middle area between 50 and 150 lp/mm shows significant variance in the MTF between microcameras, mostly due to micro optics.

Birefringence in the microcamera optics is a primary factor in degrading optical performance of the AWARE-2 camera. This birefringence was introduced into the polycarbonate element during element molding. Supplementary Figure 4.7 shows the polycarbonate lens element on the left between crossed polarizers, which is compared to the a test element molded from the low birefringence plastic material Zeonex E48R. Both the choice of material and the hexagonal flange contributed to excessive stresses and therefore birefringence in the elements. This birefringence is indicative of internal refractive index variations that produce aberrations as well as errors introduced by the polarization scrambling of the incoming rays as they propagate through the spatially varying birefringence axis medium. While the effects of the element birefringence are difficult to quantify, the performance of glass element microcameras discussed below supports the assertion that the birefringence is a major contributor to performance degradation, and that newer plastics such as the Zeonex and other optical plastics designed for lower birefringence are likely to remedy this deficiency.

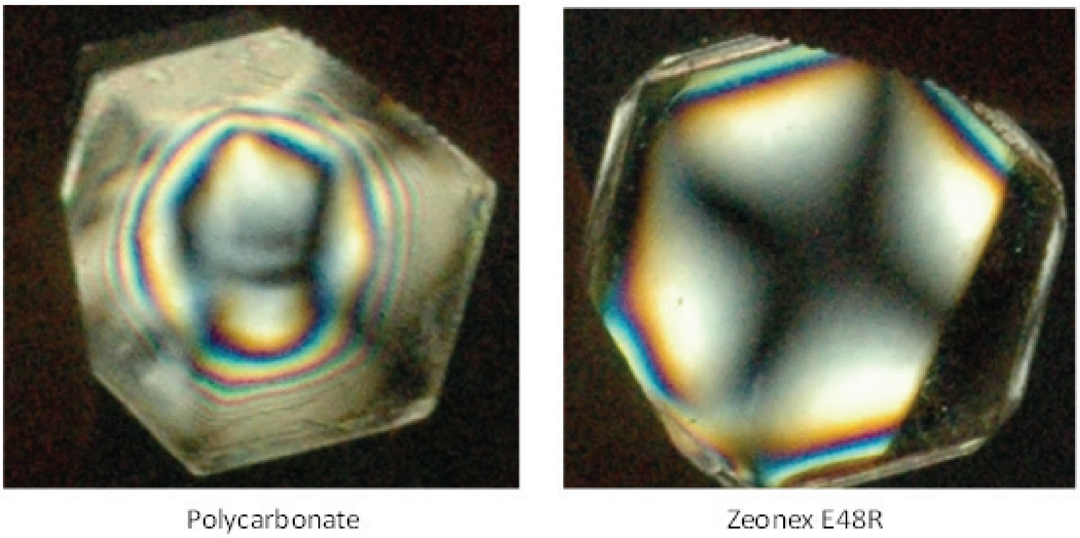

Supplementary Figure 4.7. Birefringence in the plastic microlens elements, as viewed between crossed polarizers. The polycarbonate lens exhibits many waves of retardance between the two incident polarizations. The Zeonex E48R lens exhibits far less birefringence, as evidenced by the “cross” pattern showing a fraction of a wave of birefringence in the interior of the lens, with the birefringence being radially oriented so that it is mainly visible when the extraordinary axis is not aligned with the axes of the crossed polarizers.

Supplementary Figure 4.8 explores the effects of polarization and spectral bandpass on imaging performance for a single microcamera. The bandpass filters and polarizer were placed in front of the micro optic, significantly improving the high-frequency content of the image. A nearly 10% increase in MTF at 170 lp/mm is observed for linear polarization and 550nm bandpass filter. A linear polarizer in front of the micro optic does not eliminate birefringence entirely, but only partially for the first plastic element. Additional birefringence can be seen by looking at numbers on the resolution chart. Double images can be clearly seen even in the polarized images.

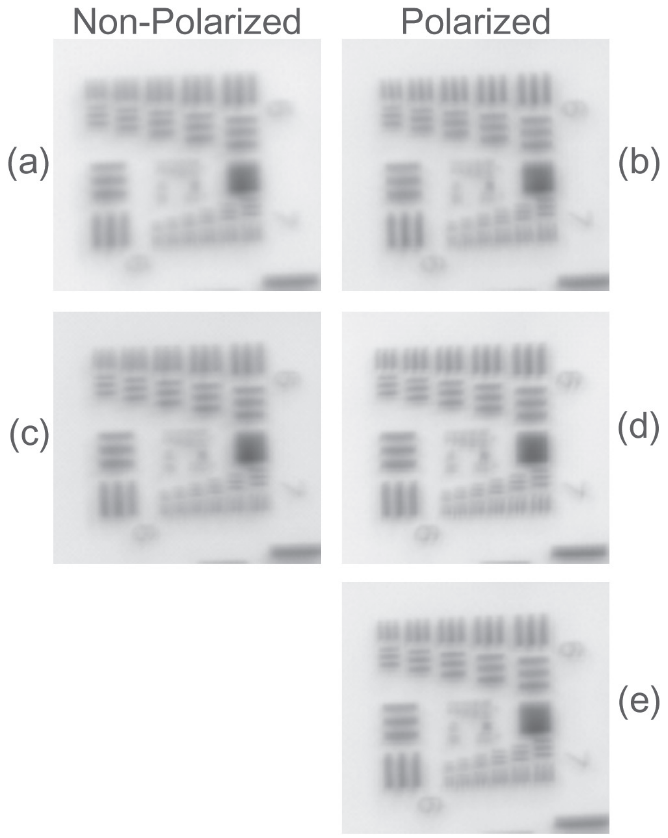

Supplementary Figure 4.8. Micro optics aberrations and effects of linear polarization. (a) non-polarized, white light; (b) linearly polarized white light; (c) non-polarized, 550nm +/- 5nm bandpass; (d) Linearly polarized, 550nm +/- 5nm bandpass; (e) 450-650nm bandpass, polarized. Note that bandpass filtering makes little difference, but polarization makes a substantial difference, indicating birefringence in the plastic optics.

The plastic microcamera optics were also tested independent from the gigagon. Supplementary Figure 4.9 shows the MTF for the plastic microcamera acting as an independent relay imager. Three different illumination sources were used, the first was white light (sunlight emulation Solux lamps), second was 550nm light with a 10nm bandpass, and lastly 450-650nm bandpass white light. The microcamera is designed to correct for chromatic aberrations of the gigagon, so using white light with just the microcamera can introduce additional aberrations. Using monochromatic light and re-adjusting for chromatic focal shift, one observes a high-frequency improvement over white light. The main point from Supplementary figure 4.9 is the significant improvement in MTF for linearly polarized light, indicating that birefringence is lowering the MTF.

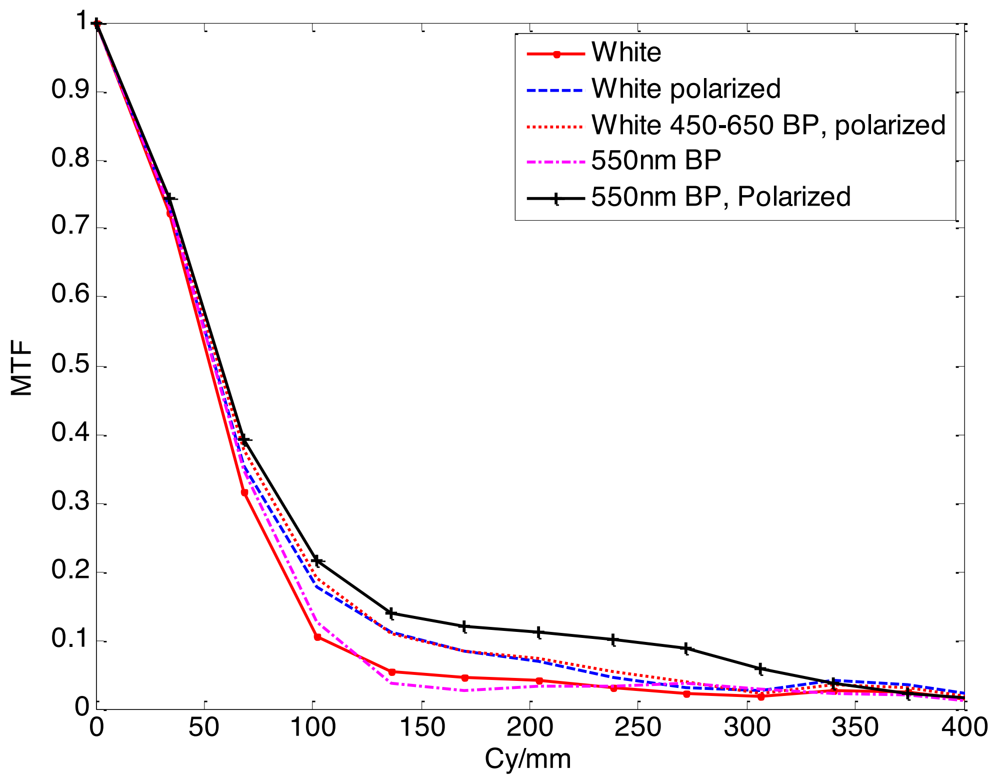

Supplementary Figure 4.9. MTF comparison for plastic micro optics. Polarization and bandpass filters increase the high-frequency content of the image significantly, but little change is apparent below 50 lp/mm.

C. Glass Microcameras

To further analyze the impact of defects in the plastic optics, an all glass micro optic barrel was designed and built. Supplementary figure 4.10 is single microcamera image taken through the gigagon using the glass optics, with zoomed in regions of the image to highlight resolution. Supplementary figure 4.11 shows a composite taken through the gigagon using three glass microcameras. The images were stitched flat-field correction followed by rotation and translation to match features in Adobe Photoshop. Overlap was blended by using the eraser tool with feathering. No distortion correction or geometric warping was done to the images to correct for perspective. More sophisticated stitching could be done, but this shows the quality of the images and how little manipulation is needed for building a high-resolution composite.

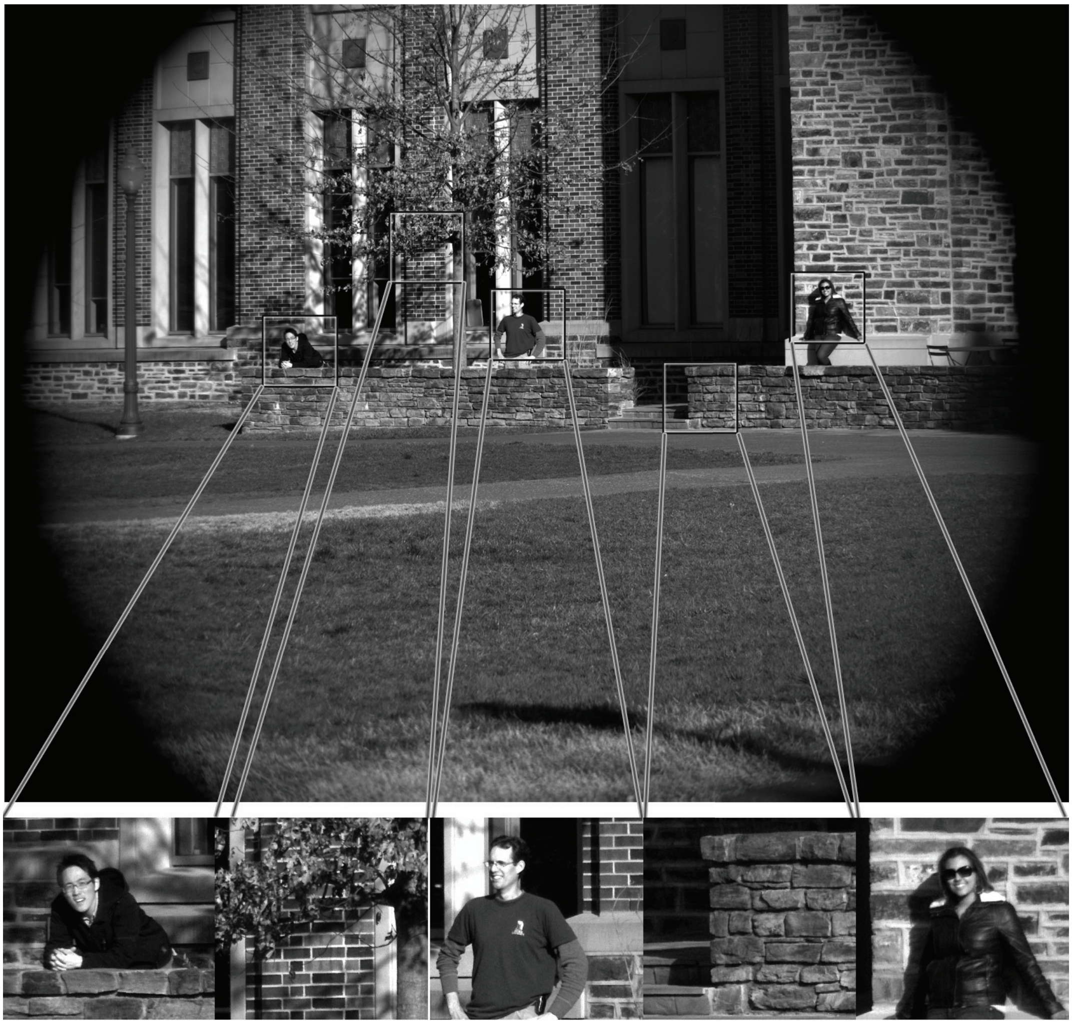

Supplementary Figure 4.10. Single microcamera image taken with glass micro optic and gigagon. The inset images are 300 pixels across.

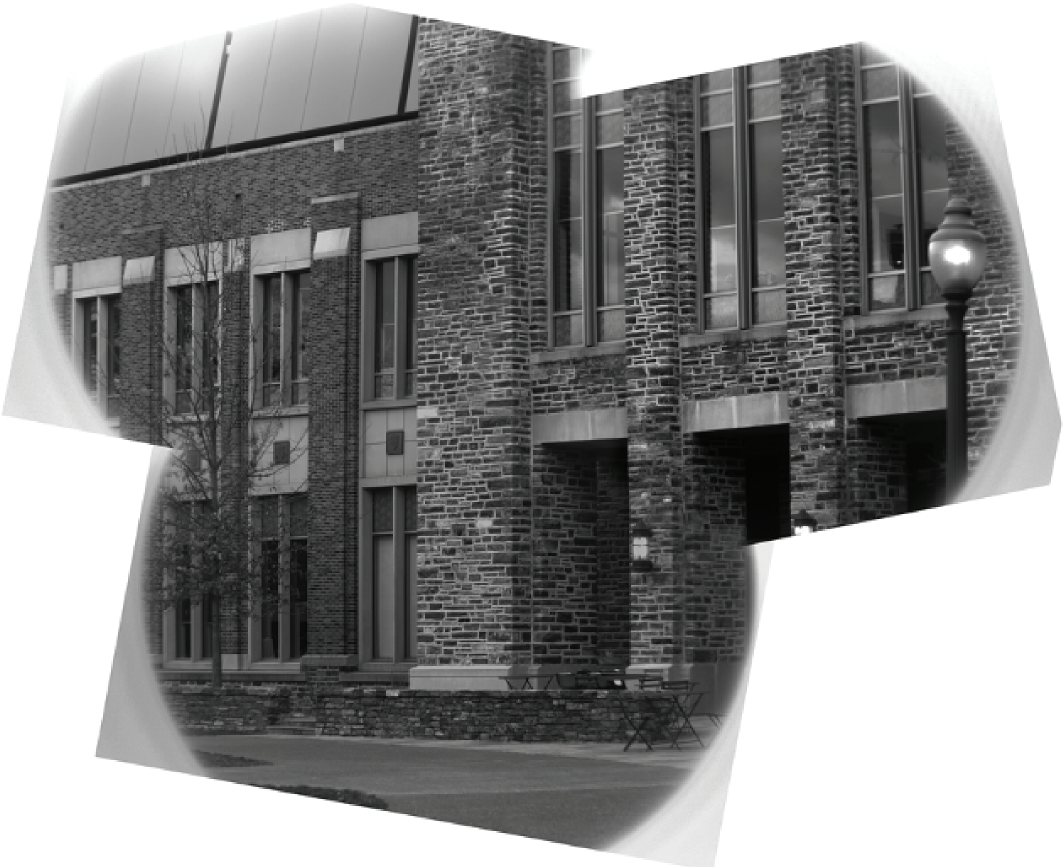

Supplementary Figure 4.11. Photoshop hand-composite using three glass micro optics in the AWARE-2 camera in hole numbers 71, 72, and 100 (refer to Supplementary Figure 4.1).

Supplementary Figure 4.12 compares a resolution target imaged using the plastic and glass micro optics alone and Supplementary Figure 4.13 shows the same comparison including the gigagon and micro optics for a large target at a significant distance.

a b

Supplementary Figure 4.12. Raw image comparisons between plastic (a) and glass (b) optics through just the micro-optics. The magnification of the two different micro optics designs is not equal, therefore the plastic optic is magnified compared to the glass optic. The barrel optics are not ultimately designed for imaging in this way, since aberrations are actually introduced into the image for correcting the gigagon. However, the imaging performance of the glass optic is pixel limited, shown in these 300% zoomed, cropped images of a chrome on glass Edmund Optics USAF target.

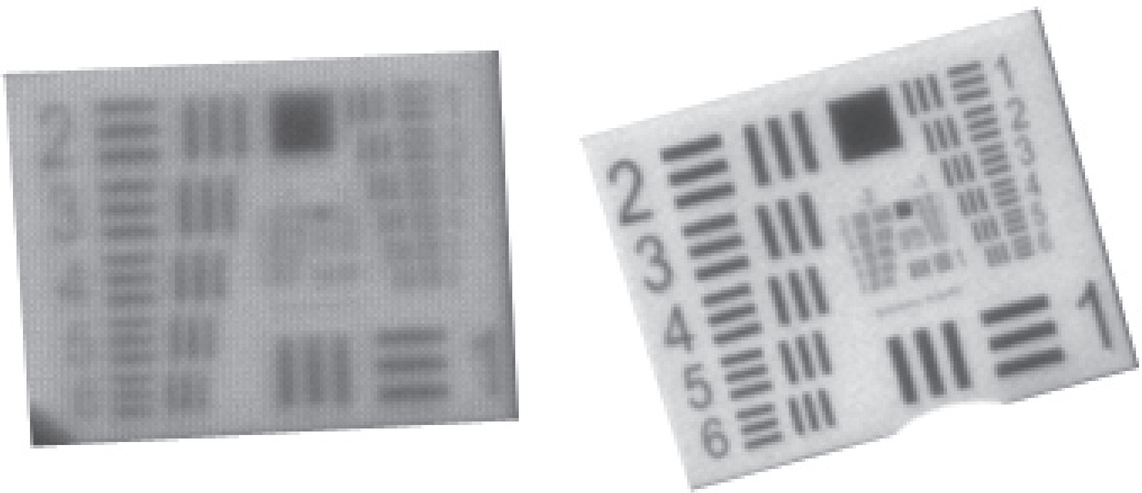

Supplementary Figure 4.13. Raw image comparisons between glass (right) and plastic (left) optics through the entire AWARE-2 camera at an AirForce resolution target. The images are shown at 200%, showing pixel limited optical resolution in the glass optics. The object distance was different for each object to account for different magnifications.

Supplementary Figure 4.14 compares the measured MTF for the plastic and glass systems. We are investigating the reasons for the discrepancies, since the model MTFs are similar. The MTF for plastic optics is shown at three different field points of the microcamera, where the camera was optimally focused for each field point. The glass microcamera shows MTF degradation at the edge of the field, but in general, this edge will rarely be used, since this extreme edge will generally be included in the adjacent microcamera. The model MTF is compensated for the pixel transfer function (PTF) of the Aptina sensor's rectangular pixels. The MTF of the Aptina sensor is assumed to be 2/pi at Nyquest, or 357 lp/mm.

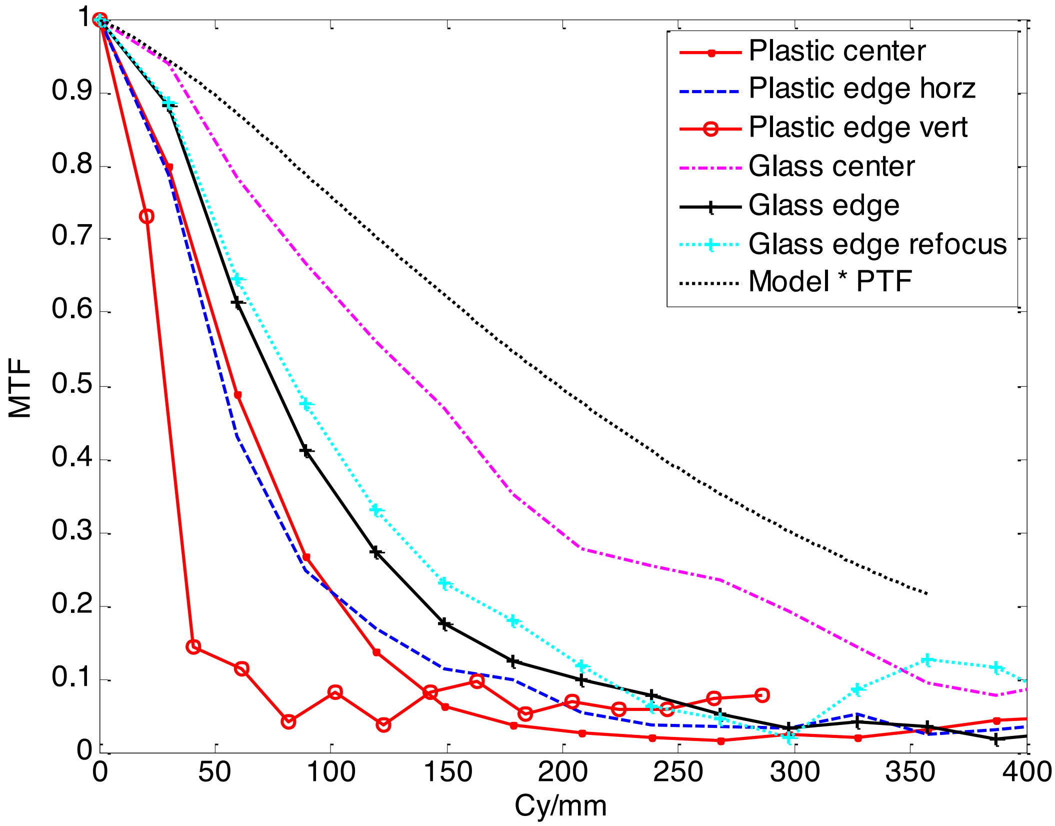

Supplementary Figure 4.14. MTF comparisons between glass and plastic micro optics and the design model which includes the pixel transfer function (PTF) for the Aptina sensor.

References

1 Marks, D. L., Tremblay, E. J., Ford, J. E. & Brady, D. J. Microcamera aperture scale in monocentric gigapixel cameras. *Appl. Opt.* **50**, 5824-5833 (2011).

2 Pollock, D. B. *et al.* Multi-lens array system and method. US Patent Application 11/456,745, (2006).

3 *Gigapixl project website*, < http://www.gigapxl.org/>

4 Leininger, B. *et al.*Autonomous real-time ground ubiquitous surveillance-imaging system (ARGUS-IS). 69810H-69811 (SPIE, 2008).

5 *LSST Overview document.*, < http://arxiv.org/abs/0805.2366> (
6 *Pan-STARRS website.* , <http://pan-starrs.ifa.hawaii.edu/public/home.html>
7 Son, H. S., Marks, D. L., Hahn, J., Kim, J. & Brady, D. J. Design of a spherical focal surface using close-packed relay optics. *Opt. Express* **19**, 16132-16138 (2011).
8 Dean, J. & Ghemawat, S. MapReduce: simplified data processing on large clusters. *Commun. ACM* **51**, 107-113, doi:10.1145/1327452.1327492 (2008).
9 Lowe, D. G.Object recognition from local scale-invariant features. in *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on.* 1150-1157 vol.1152.
10 Bay, H., Ess, A., Tuytelaars, T. & Van Gool, L. Speeded-Up Robust Features (SURF). *Computer Vision and Image Understanding* **110**, 346-359, doi:10.1016/j.cviu.2007.09.014 (2008).