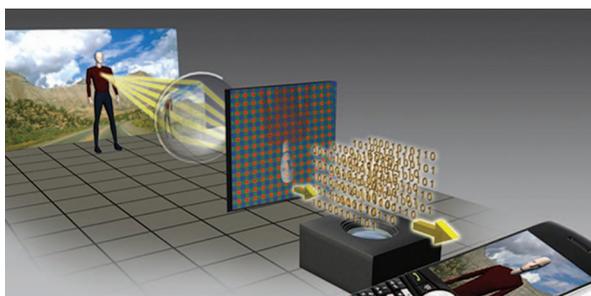


# A Vision for the Future

By David Moloney and Oscar Deniz Suarez

For the past 40 years, computer scientists and engineers have been building technology that has allowed machine vision to be used in high-value applications from factory automation to Mars rovers. However, until now, the availability of computational power has limited the application of these technologies to niches with a strong enough need to overcome the cost and power hurdles. This is changing rapidly as the computational means have now become available to bring computer vision to mass-market applications in mobile phones, tablets, wearables, drones, and robots, enabling brand-new user experiences within the cost, power, and volumetric constraints of mobile platforms.

Computational imaging represents a historic transition from the 150-year-old paradigm of taking photos on silver halide photo stock and chemically developing and printing them to computing (making) pictures. According to Hayes [1], a digital camera is an image-creation device rather than a simple, passive recording device. In the existing digital still cameras, the focus is on making digital images identical to their chemical forebears. However, once the camera contains sufficient image-processing horsepower, a computational camera can move beyond the reality captured by conventional digital cameras. The sensor



**FIGURE 1.** The imaging and vision capture and processing chain.

array in such a camera fills the role previously played by film, but it is the beginning rather than the end of the image creation process.

In his book describing the work of photography pioneer Ansel Adams in the 1920s, 1930s, and beyond, William Turnage [2] claims Adams spent up to a day per print effectively doing manually what today would be called high-dynamic-range (HDR) photography: “[Adams] always said that the negative is the equivalent of the composer’s score and the print is the equivalent of the conductor’s performance.” Similarly, the concept of multiaperture (array) and light field (plenoptic) cameras, which allow depth to be recovered from images and enable a range of depth of field and other effects to be implemented computationally, have been around since the turn of the 20th century, since Lippmann demonstrated his  $3 \times 4$  array camera in 1911 [3], essentially waiting 100 years for the computational means to catch up with his vision.

The current computational photography architectures [4]–[6] focus very much on image capture and postprocessing in

the point-and-shoot camera paradigm within existing smartphone architectures. In this model, image capture is controlled using application programming interfaces (APIs) such as fCam, Android Camera 2.0, and the forthcoming Camera 3.0 API for Android (derived from fCam). These APIs are aimed at opening up the previously closed world of the camera

and, to some extent, the camera image signal processing (ISP) internals to the application developer (Figure 1).

This fine control of the capture process enables applications based on focal stacks (rapidly captured bursts of images with different camera settings) in smartphones such as HDR image and video capture, synthetic aperture photography where the depth of field can be varied to focus on different parts of the image, removal of unwanted persons or objects within images, best shots of groups (everybody smiling, facing the camera and not blinking, etc.), and sophisticated ISP functions such as alignment of multiple raw images for super-resolution.

The Android Camera 3.0 API makes extensive use of the burst mode, capturing sequences of high-resolution images, and associated depth information, at high rates in smartphones, leveraging existing GPU cores using computational imaging algorithms that use RenderScript to deliver the best images in a platform-independent manner. Similar efforts to standardize camera APIs, the capture and fusion of sensor metadata,

and hardware acceleration for computer vision are ongoing within the Khronos industry-driven standards organization as well as more domain-specific initiatives such as advanced driver assistance systems within the automotive community driven by EU NCAP and U.S. safercar safety standards.

The APIs and programming models proposed in [4]–[6] aim to leverage the existing hardware processing resources by layering software to marshal the heterogeneous computational resources in application processors (APs) to do a new job for which they were never designed. Generally, the hardware architectures underlying these APIs, and currently used to deliver these advanced camera functions, are multicore CPUs, DSPs, and GPUs already present in existing or derivative APs, which are widely used in smartphones and tablets. The result is a compromise that constrains the computational capability of the system and delivers poor user experiences and very poor battery life.

## MOVING BEYOND TAKING PICTURES

One of the defining characteristics of computer vision is that it turns into a means of making measurements and inferences about the world and taking action based on those measurements rather than simply capturing a scene by measuring the light incident at each pixel in an array. On a very basic level, many of us have unwittingly been using massive numbers of cameras to measure the world for the past 15 years since Microsoft introduced the first camera mouse in 1999. These optical mice use a combination of an LED mounted at a glancing angle and a simple camera to extract the underlying texture in any surface the mouse is moved across, with postprocessing to extract a motion vector, which is used to control the cursor on a PC screen. In the past few years, new human interaction devices such as Microsoft's Kinect, Leap Motion's camera-based gesture device, and Tobii's eye-tracker have begun to revolutionize the way we interact with digital content on PCs and gaming consoles.

More recently, devices such as HP's Sprout (Figure 2) use a combination of



FIGURE 2. The HP Sprout workstation.

a physical LCD touchscreen and a downward-facing video projector coupled with a computer-vision-enabled horizontal touch surface (TouchMat) to build a user interface that extends outside the box, allowing the user to interact with the real world in new and exciting ways, capturing and manipulating 2-D and 3-D digital content in a highly intuitive manner. These capabilities have also begun to appear in mobile devices such as Amazon's FirePhone [26], which boasts four dedicated computer-vision (global-shutter) cameras mounted in the corners of the device along with associated infrared LEDs in addition to the conventional front- and rear-facing cameras. The four cameras allow novel UI features such as viewpoint-dependent 3-D rendering, which adapts to the user's head position, HDR image capture, and cloud-backed image search.

The next step is, of course, to move beyond the limitations of our personal devices entirely using telepresence and autonomous drones and robots as well as embedding vision in a broader range of products. A key enabler for these use cases is simultaneous localization and mapping (SLAM), widely used by autonomous robots operating in unknown environments and developed in the early 1990s for Mars rovers [7], [8]. In such systems, SLAM software turns a conventional RGB camera into a six-degree-of-freedom (DoF) transducer, and the autonomous mobile platforms enabled by it must be able to reliably measure ego-motion as otherwise long-term navigation is impossible, and conventional means of location determination versus a known reference such as GPS are unavailable. Stereo odometry determines the ego-motion of a stereo camera in the 6 DoFs

that are possible in the 3-D world (three for translation and three for rotation) and compensates for wheel slippage, which can otherwise cause problems with odometry. Furthermore, sensor fusion and tracking are also integral components of many autonomous vehicles and robots.

Until very recently, SLAM algorithms have been academic in nature and have not been optimized for embedded platforms. An approach to implementing SLAM on embedded platforms leveraging SIMD coprocessors, DSP, and multicore CPUs is outlined in [9], and rapid progress in embedding such technology is being made through initiatives such as Google's Project Tango (Figure 3) [19] and SLAMbench [10]. A good example of a consumer device incorporating SLAM is the Dyson 360 Eye robotic vacuum cleaner (Figure 4). [25], and we can expect that such systems will rapidly fall in price and increase in capabilities in the coming years driven by advanced semiconductors.

Returning to human vision, the idea of virtual reality has been around for a long time, but the experience never delivered on the promise with many of the devices either not working or inducing nausea in the unfortunate wearer. The key issue according to Abrash [11] is latency or, in other words, the delay between head motion and the corresponding virtual world update reaching the eyes. Too much latency results in images drawn in the right place but at the wrong time, creating anomalies that are amplified by



FIGURE 3. The Project Tango phone.



FIGURE 4. The Dyson 360 Eye.



**FIGURE 5.** The rendering of mixed reality by Magic Leap.

head motion; the faster the head moves, the greater the anomaly, and it is compounded further when the head changes direction. Moving beyond the confines of traditional VR devices such as Oculus Rift, companies like Magic Leap [24], who recently raised more than US\$542 million from Google and other investors, hint at an intriguing world of seamless mixed reality, blending rendered graphics with reality using advanced displays and enabled by computer vision (Figure 5).

Latency is even more demanding in applications such as active headlights in cars; an example is the Carnegie Mellon University Smart Headlight program directed by Prof. Takeo Kanade [12], where the total round-trip latency from snowflake or raindrop reflection through sensor computer-vision hardware and closing the loop by modulating an LED projector array in the headlight is on the order of 1,500 ms.

Another important application for computer vision is in giving an out-of-body-experience by mounting a camera on a drone and having it track the user, for instance, to make sports and activity videos. These devices have just begun to appear on the market, and it can be expected that such devices will proliferate the consumer electronics space in the coming years as prices fall and fly time increases. The weight of both batteries,



**FIGURE 6.** The ProxDynamics PB100 drone.

the electronics, the airframe, and the sensor array are key concerns with the electronics often being the limiting part of programs such as sFly [22], which achieved a 10-min autonomous fly time and required about 1 W for a 100-g payload. The military, law-enforcement, and emergency services are being targeted by dedicated price-is-no-object devices such as the ProxDynamics Black Hornet PB100 drone helicopter (Figure 6) [23], which weighs in at a miniscule 16 g (equivalent to three sheets of A4 paper) and boasts three cameras and a wireless link that operates to over 1 km in range.

Beyond UI, smart cameras can also be used to measure other phenomena; for instance, Eulerian Video Magnification [13] can be used to amplify natural

Another important application for computer vision is in giving an out-of-body-experience by mounting a camera on a drone and having it track the user.

movement, luminance, or color changes and make it visible to the human eye much in the same way as time-lapse photography allows us to see plants grow or clouds move across the sky. We can expect such features to appear in medical equipment and baby monitors in the not-too-distant future. Eulerian magnification can also be used to reconstruct sound [14] that causes a plant's leaves to move in a sealed room when fed into an appropriate inverse acoustical model; however, the civilian applications for this technology are less clear.

Computer vision and computational photography are undoubtedly in a rapid growth phase with new vision-based applications appearing on a regular basis. One prominent example is Placemeter, which uses public video feeds (from old unused smartphones attached to windows) and computer vision algorithms to create the first ever real-time layer of data about places, streets, and neighborhoods.

Placemeter [21] collects and serves up-to-the-minute information like how crowded a place is, how long the wait is, and whether it will get more or less crowded in the next hour. There are many more examples, including image search, panoramas, face detection in smartphones' cameras, face recognition biometrics to unlock devices, video stabilization in YouTube, and Facebook's facial recognition for photo tagging.

## EYES EVERYWHERE

It is clear that computer vision is in an explosive growth phase transitioning from traditional automation in factories to the world in which people work, live, and play. While computer vision is a mature research field from a theoretical point of view, practical ubiquitous vision has not progressed to the same extent because the systems have been confined to labs and factories by the lack of suitably power-efficient hardware on which to run mobile applications. The underlying difficulties are primarily the computational and cost requirements imposed by consumer devices. Human vision is immediate; we open our eyes and can immediately recognize and categorize objects and the structures of scenes. What we do not realize is the vast computational resources that our brain brings to bear unconsciously to make all of this happen. When our eyes are open, vision accounts for two-thirds of the electrical activity of the brain, and the sensors, our eyes, are located right there in the same casing as our brains. The fact that our brain manages to deliver all of this functionality, which we can only dream of duplicating in a machine, is even more amazing when we realize that our brain consumes a mere 20 W and is fueled by environmentally friendly renewable sugars. This being said, even our crude approximations of human vision are now starting to yield useful results in practical power and cost envelopes.

With the advent of cloud computing, it is tempting to think that the cloud alone can bear the computational load associated with vision processing. However, crunching the enormous volume of

visual data is currently beyond the reach of all but a few very large companies, and network bandwidths cannot cope with the massive use of cameras. Power efficiency is a major issue, and wirelessly transmitting data for remote computation can cost up to one million times more energy per operation compared to processing locally in a device.

A further reason for pushing processing and even limited decision making to the network edge in applications such as virtual or augmented reality, autonomous vehicles, and robots is that data centers have major issues with latency. Finally, privacy in the cloud is a concern, especially when we consider that more often than not, the subject of our monitoring will be humans. Our own brains are colocated with our eyes, processing visual data close to the point of origin, and our reflexes deal with the latency of transmitting stimuli to our brain and back to the muscles using local neural circuits to save time and get us out of “harm’s way.”

Placing the computational resources close to optical and other sensors in cyber-physical systems clearly solves many of the same issues. Going beyond this, perhaps we require not only “intelligence everywhere” but also “eyes everywhere” for many applications. In the scientific and technical literature, the closest systems to this eyes everywhere paradigm are to be found under the terms “embedded vision” and “vision sensor networks.” On one hand, embedded vision refers to vision systems that are integrated into more complex devices such as automotive or medical equipment. Those systems are a natural evolution of fixed industrial vision systems based on PCs and smart cameras. While such systems are increasingly being used, the power and size requirements are not so stringent and development is typically application specific. Vision sensor networks, on the contrary, aim at smaller, power-efficient vision sensor nodes that can operate as stand-alone nodes; however, the proliferation of standards and the lack or interoperability has frustrated widespread adoption and commercialization, causing some researchers [15]

to observe: “Given the fair number of proposed designs, it is somewhat surprising that a general-purpose embedded smart camera based on an open architecture is difficult to find and even more difficult to buy.”

The current heterogeneous platforms such as mobile phone APs can be used to prototype algorithms, but, from the published results, they appear limited to VGA 30-frames/s resolution and around 5-W power dissipation. Technologists such as John Carmack, the chief technology officer of Oculus [17], say that AP power dissipation limits realistic VR experiences to around 15 min. Even then, implementing relatively simple



**Movidius is enabling a swathe of new computer vision applications to be brought to the mass market for the very first time in embedded devices such as mobile phones, tablets, and cameras.**

pipelines on these platforms involves a high level of complexity involved in the coordination of multiple ARM processors, NEON SIMD extensions, and GPUs, all of which must share access to data structures in memory through shared access or copying data. In the case where data are copied to and from the GPU, it means that a certain level of granularity in terms of GPU tasks is required to make the effort of moving data to and from worthwhile. This complexity must be managed explicitly by the programmer, with a mixture of C/C++ code for the ARM, SIMD assembler for NEON, and OpenGL ES or OpenCL shaders, or via vendor-supplied libraries and APIs.

Other options, such as miniature IP and smart cameras, do not come close to the required size, energy consumption, cost, or processing power. On a related note, in 2014, Freescale released the wearable reference platform (WaRP) [16]. The WaRP is the first attempt to provide a reference platform for the

future development of wearable devices. While it is open, small, and can be connected to cameras, the WaRP has not been designed for mobile-embedded vision, which is the most challenging capability in terms of required processing power and energy consumption.

In this context, innovative vision applications are typically based on (little-effective) DIY kits and smartphones or supported by large companies that can fund the specific hardware and software developments needed. Research and academia do not seem to have a versatile platform for deploying innovative vision applications and for rapidly designing new products based on the latest advances in the field.

## **MYRIAD OF NEW POSSIBILITIES**

In the aforementioned context, Movidius is focused on bringing human vision and scene understanding to mobile devices, allowing low power, always-on vision capabilities in devices with the very low latencies required for interactive services, self-driving cars, and robots.

Movidius is enabling a swathe of new computer vision applications to be brought to the mass market for the very first time in embedded devices such as mobile phones, tablets, and cameras. The first generation of the Movidius Myriad Vision Processing Unit (Myriad1 VPU) powers the computer vision subsystem in Google’s Project Tango (see Figure 3), where it handles all of the high-performance ISP, feature tracking, and tracking tasks using ten times less power than any other solutions available on the market today [19]. In August 2014, Movidius introduced its second-generation Myriad2, which was designed to achieve 20–30 times the processing per watt of the previous generation, Myriad1. Myriad2 is a system-on-chip that embeds a software-controlled multicore, a multiported memory subsystem, and caches that can be configured to allow a large range of workloads to be handled, providing exceptionally high sustainable on-chip data and instruction bandwidth to support the 12 processors, two RISC processors, and high-performance video hardware accelerator filters. Supporting

a sustained throughput of 600 mp/s means that Myriad2 can deliver 1,080p120 with less than 20% of the available pixel bandwidth. As a result of the highly power-efficient architecture of Myriad2, an OpenCV compatible multiscale Haar Cascade consisting of 20 stages, computed using 12 SHAVEs and one of the HW accelerators in Myriad2, can calculate 50,000 multiscale classifications for each 1,080-p resolution frame in less than 7 ms (a key requirement for immersive VR [11]).

Because of this performance and the focus on embedded vision systems, Myriad2 has been recently selected to power the “Eyes of Things” (EoT) platform [18], an innovation action funded by the European Union’s Horizon 2020 Framework Programme for Research and Innovation. The objective is to build an embeddable always-on computer vision system that can be used as a generic platform for any applications requiring mobile/embedded vision. The hardware and software infrastructure developed in the EoT will be available to OEMs and the general public, and the applications for EoT nodes are myriad, including wearables, UAVs, robotics, and surveillance. It is estimated that the use of the EoT platform will save up to 41% in the time to market for advanced vision-based applications.

## CONCLUSION

It is clear that a wide range of computational photography and computer vision techniques can greatly enhance user experiences, and, while these techniques are well understood, they are in a considerable state of flux, which precludes the use of fixed-function hardware as there is no one-size-fits-all device or algorithm for many or any of these applications. In this context, a software-programmable device that is optimized for vision workloads clearly offers an optimal balance of power and performance. It is our belief that VPU will become the industry standard way of handling such workloads using open APIs such as Khronos OpenVX [20] in the same way that GPUs are used to process computer graphics workloads, and Myriad is simply the first example of this class of

device to make it to the market. As we see it, the availability of capable, flexible, and power-efficient homogeneous processor architectures created for computer vision will allow users to begin to expect to “live in the future” where always-on computer vision applications are possible at HD resolution, dissipating only hundreds of milliwatts. These platforms will be utilized in a range of use cases, including stand-alone AR displays, wearable devices, UAVs, mobile robots, surveillance cameras, and tablets and phones where the computer vision processor offloads demanding tasks from the AP. For the first time, the device designers will focus on designing world-beating products with revolutionary CV functionality in record time scales and with exceptionally low power requirements.

## ABOUT THE AUTHORS

**David Moloney** (david.moloney@movidius.com) earned his B.Eng. degree in electronic engineering at Dublin City University in 1985 and his Ph.D. degree from Trinity College Dublin in 2010 in the area of FPGA-based HPC for computational fluid dynamics. He is the chief technology officer of Movidius and has worked in the semiconductor industry internationally for the past 28 years with Infineon in Germany, ST Microelectronics Italy, Parthus-Ceva (Ceva-DSP), and FrontierSilicon in Ireland. He founded Movidius in 2005 with Sean Mitchell. He has 31 granted patents and numerous conference and journal publications. He is a reviewer for *IEEE Communications Magazine*, a steering committee member for the U.K. EPSRC PAMELA project, and an expert reviewer for the European Union (EU) Commission on programs such as ARTEMIS. He is a member of the EU FP7 HiPEACNoE and collaborates on the FP7 PHEPPER and EXCESS projects as well as the EU Horizon 2020 Eyes of Things (EoT) project. His interests include processor architecture, computer vision algorithms and HW acceleration and systems, hardware and multiprocessor design for DSP communications, HPC, and multimedia applications.

**Oscar Deniz Suarez** (oscar.deniz@uclm.es) earned his M.Sc. degree in computer science and his Ph.D. degree in

computer vision from the University of Castilla-La Mancha. He is an associate professor at the University of Castilla-La Mancha. He contributes to the VISILAB research group his interests including computer vision and machine learning. He is the author of more than 50 refereed papers in journals and conferences. He received the runner-up award for the best Ph.D. work on computer vision and pattern recognition by the Spanish Association on Pattern Recognition and the “Image File and Reformatting Software” Challenge Award by Innocentive, Inc. He was a national finalist of the 2009 Corbaayen Award. He has led projects for more than ten cutting-edge companies in the past three years, most of them for mobile vision application development. He has two patents. He has also contributed to OpenCV, the well-known open-source computer vision library. He is the author of three books on OpenCV and OpenCV programming for mobile devices. He has also served as visiting researcher at Carnegie Mellon University, Pittsburgh, Pennsylvania, Imperial College London, and Leica Biosystems, Ireland. He is a Senior Member of the IEEE and is affiliated with the AAI, IAPR, and The Computer Vision Foundation. He serves as an academic editor of *PLoS ONE*. Currently, he is the coordinator of the EU Horizon 2020 Eyes of Things project and a partner in the FP7 AID-PATH Marie Curie Action.

## REFERENCES

- [1] B. Hayes, “Computational photography,” *Amer. Scientist*, vol. 96, no. 2, pp. 94–99, Mar.–Apr. 2008.
- [2] W. A. Turnage, *Ansel Adams: Our National Parks*, 1st ed. New York: Little, Brown and Company, May 21, 1992.
- [3] “Integral photography: A new discovery by Professor Lippmann,” *Scientific Amer.*, vol. 105, pp. 164, Aug. 1911.
- [4] S. H. Wang and U. Timofei. (2013). Computational photography. A new way to look at image capture. [Online]. SAMSUNG Whitepaper. Available: [http://www.samsung.com/global/business/semiconductor/minisite/Exynos/data/computationalphotography\\_whitepaper.pdf](http://www.samsung.com/global/business/semiconductor/minisite/Exynos/data/computationalphotography_whitepaper.pdf)
- [5] NVIDIA. (2013). Chimera™: The NVIDIA computational photography architecture. [Online]. Available: <http://www.nvidia.com/object/white-papers.html>
- [6] Qualcomm. (2014, Feb. 5). Qualcomm Snapdragon 800 product brief. [Online]. Available:

<https://www.qualcomm.com/media/documents/files/snapdragon-800-processor-product-brief.pdf>

[7] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*, I. J. Cox and G. T. Wilfong, Eds. New York, NY: Springer-Verlag, 1990, pp. 167–193.

[8] J. J. Leonard and H. F. Durrant-Whyte, "Mobile robot localization by tracking geometric beacons," *IEEE Trans. Robot. Automat.*, vol. 7, no. 3, pp. 376–382, 1991.

[9] B. Vincke, A. Elouardi, and A. Lambert, "Real time simultaneous localization and mapping: Towards lowcost multiprocessor embedded systems," *EURASIP J. Embed. Syst.*, vol. 5, no. 1, pp. 1–14, 2012.

[10] L. Nardi, B. Bodin, M. Z. Zia, J. Mawer, A. Nisbet, P. H. J. Kelly, A. J. Davison, M. Luján, M. F. P. O'Boyle, G. Riley, N. Topham, and S. Furber. (2014, Oct. 8). Introducing SLAM-Bench, a performance and accuracy benchmarking methodology for SLAM. [Online]. Available: <http://arxiv.org/abs/1410.2167v1>

[11] M. Abrash. (2012, Dec.). Latency—The sine qua non of AR and VR. [Online]. Available: <http://blogs.valvesoftware.com/abrash/latency-the-sine-qua-non-of-ar-and-vr/>

[12] R. Tamburo, E. Nurvitadhi, A. Chugh, M. Chen, A. Rowe, T. Kanade, and S. G. Narasimhan, "Programmable automotive headlights,"

in *Proc. European Conf. Computer Vision (ECCV)*, Lecture Notes in Computer Science, Vol. 8692, 2014, pp. 750–765.

[13] H. Y. Wu, M. Rubinstein, E. Shih, J. Guttag, F. Durand, and W. T. Freeman, "Eulerian video magnification for revealing subtle changes in the world," *ACM Trans. Graph. (Proc. SIGGRAPH 2012)*, vol. 31, no. 4, pp. 65:1–65:8, 2012.

[14] A. Davis, M. Rubinstein, N. Wadhwa, G. Mysore, F. Durand, and W. T. Freeman, "The visual microphone: Passive recovery of sound from video," *ACM Trans. Graph. (Proc. SIGGRAPH 2014)*, vol. 33, no. 4, pp. 79:1–79:10, 2014.

[15] B. Murovec, J. Pers, R. Mandeljc, V. Kenk, and S. Kovacic, "Towards commoditized smart-camera design," *J. Syst. Architect.*, vol. 59, no. 10, pp. 847–858, 2013.

[16] Freescale. (2015, Feb. 5). Wearable reference platform. [Online]. Available: <http://www.warpboard.org/>

[17] J. Gaudiosi. (2015, Feb. 5). Oculus VR CTO John Carmack: The extended Q&A. [Online]. Available: <http://fortune.com/2014/12/29/oculus-vr-john-carmack-extended/>

[18] (2015, Feb. 5). [Online]. Available: <http://eyesofthings.eu>

[19] D. Cardina. (2015, Feb. 5). Movidius, the chip maker behind Google's Tango, wants to be the king of computational photography. [Online]. Available: <http://www.extremetech.com/>

extreme/187232-movidius-the-chip-maker-behindgoogles-tango-wants-to-be-the-king-of-computational-photography

[20] Khronos Group. (2015, Feb. 5). Portable, power-efficient vision processing. [Online]. Available: <https://www.khronos.org/opencv/>

[21] Placemeter Inc. (2015, Feb. 5). Placemeter. [Online]. Available: <http://www.placemeter.com/>

[22] (2015, Feb. 5). Swarm of micro flying robots. [Online]. Available: <http://www.sfly.org/>

[23] A. Tarantola. (2015, Feb. 5). Black Hornet: The \$195,000 spy plane that fits in the palm of your hand. [Online]. Available: <http://gizmodo.com/5981975/black-hornet-the-195000-spy-plane-that-fits-in-the-palm-of-your-hand>

[24] J. D'onfro. (2015, Feb. 5). GOOGLE'S \$500 MILLION MAN: Meet the 'Weird' Guy trying to invent a new computing platform. [Online]. Available: <http://uk.businessinsider.com/rony-abovitz-magic-leap-2014-12?r=US>

[25] Dyson Inc. (2015, Feb. 5). Robot Dyson 360 Eye. [Online]. Available: <https://www.dyson360eye.com/>

[26] Amazon.com, Inc. (2015, Feb. 5). Amazon Fire Phone, 32GB (AT&T). [Online]. Available: [http://www.amazon.com/Fire\\_Phone\\_13MP-Camera\\_32GB/dp/B00EOE0WKQ](http://www.amazon.com/Fire_Phone_13MP-Camera_32GB/dp/B00EOE0WKQ)



# IEEE Access

The journal for rapid open access publishing

## Become a published author in 4 to 6 weeks.

### IEEE Access is a multidisciplinary journal that allows you to:

- Publish articles in 4 to 6 weeks
- Expect the IEEE trusted peer review with a typical **one-third** acceptance rate
- Reach millions of global users through the IEEE Xplore® digital library with free access to all
- Submit multidisciplinary articles that do not fit neatly in traditional journals
- Integrate multimedia and track usage and citation data for each published article
- Connect with readers through commenting
- Publish without a page limit for **only \$1,750** per article

IEEE Access... a multidisciplinary open access journal that's worthy of the IEEE.



Learn more at:  
[www.ieee.org/ieee-access](http://www.ieee.org/ieee-access)

**IEEE**  
Advancing Technology  
for Humanity

14-PUB-246 12/14