Alan C. Bovik

What You See Is What You Learn

I find teaching to be an extremely personal experience. After all, there are few occupations where one encloses herself or himself in a room for 50 hours or more, over a period of five months, with the same group of dozens (or more) individuals, all listening closely (we hope) as the professor pontificates upon the subject of his or her most intensive personal inquiries and study that may have spanned decades. Since I am to explain my own personal take on the experience of teaching a vital and sensorial topic, I will take the first person throughout. Hopefully, this will make it possible for me to better convey my own sense of joy and satisfaction working in the field of video processing education, as well as better justifying my views of the process of teaching this subject.

I have always viewed the subject of image and video (hereafter video; I take video to subsume image) processing as having broad appeal, owing to its relationship to such glamorous applications as cinema and television and to such recognized transformative sciences as medical imaging and astronomy. It is also accessible to our visual sensory apparatus, arguably the richest source of data that we capture during our daily experience of life. While teaching at the university level is always enjoyable, for these and many other reasons, I find teaching and learning the topic of video processing to be exceptionally fun.

In recent years, I have found that digital video processing education has become more exciting than ever because of an absolute explosion of first-of-a-kind consumer products, remarkable medical imaging modalities, amazing cinematic presentations, and a gigantic volume of images and streaming videos on the Internet and over wireless. Many of us carry digital video communication and display devices in our pockets (think of the BlackBerry and iPhone), and view digital images and videos in increasingly larger formats and higher resolutions on home entertainment centers, on our personal computers, and elsewhere. As I write this, the all-digital and immensely popular three-dimensional (3-D) movie *Avatar* has, within the past few months, revised our view of what movies are all about, both in the theatre and at home. Just a few weeks ago, the iPad was released to the public by Apple, Inc., with largely positive reviews that were also somewhat puzzled regarding just where this device, and later ones like it, might take us. While the answer to that lies in the future, I feel safe in saying that digital images and videos will become even more pervasive in our daily experience. This increasing relevance to daily life continues to drive innovations in the communications, Internet, movie production, and display industries, giving rise to significant job opportunities that require digital video processing expertise. In many senses, I think that digital video is finally arriving as a conspicuous centerpiece of modern technology. I say "finally" since this development has always seemed inevitable to me. Yet, I think that enormous amounts of innovation lie ahead, as the digital visual experience becomes richer, higher resolution, larger, immersive, and experiential.

As an educator viewing these developments, I feel that it is important to deeply consider the topic of how to best optimize the educational experience of students learning this unique topic. During the more than 25 years I having been teaching this subject in the classroom, in the laboratory, and at companies, I have come to recognize a number of elements that I feel require special emphasis. Some of these, of course, arise from my own perspectives and personal style of instruction, but others, I think, are broadly relevant. I would add that, while most of my observations are germane to the undergraduate experience, along the way I also give suggestions on improving the educational preparation of graduate students who are seeking to specialize in digital video processing.

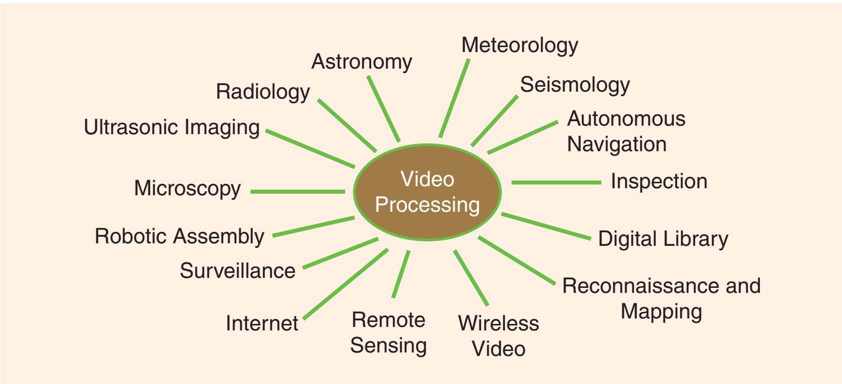

[FIG1] A part of the universe of video processing applications.

Digital Object Identifier 10.1109/MSP.2010.937497

A DIVERSITY OF APPLICATION AREAS … AND STUDENTS

I regularly teach an upper-level undergraduate course titled "Digital Image and Video Processing." One of the first things that I make clear to the class is that the field of video processing is remarkably diverse in the range of applications that it finds. Truly, there are few fields of science or engineering (if any) that do not benefit by exploiting the visual sense to better communicate and interpret data. Figure 1 depicts this diversity, and it is very far from complete. Certainly this diversity makes the field more interesting, creates numerous cross-disciplinary collaborations, and constantly opens up new avenues of practical inquiry.

However, there is another reason why I want to discuss the diversity of applications to be found in this field: a matching diversity of student interest. Because of the near-ubiquity of digital images and videos as tools, there are growing numbers of students from diverse disciplines that want to take my course. I am certain that this must be the case at other schools. I am regularly approached by students from every branch of engineering and computational science, from every branch of hard science, and even from the softer sciences, at both the undergraduate and graduate levels, regarding the feasibility of enrolling in my course. Usually, such students (even within engineering) feel trepidation about taking the course, since they naturally assume that there must be a lot of material that they will be unfamiliar with, that would be required knowledge to succeed in the course.

Long ago I made the decision that I should be receptive to, and highly encouraging of, such students. Since most other engineering and science curricula do not contain courses that teach image processing, such students must otherwise generally rely upon software library manuals, textbooks, and online references to come up to speed on the subject. This is, I feel, a questionable way to embrace the topic, and generally leads to the student, in their own work, patching together library image processing routines in a daisy-chain manner to achieve some end. Frequently, basic and often fatal errors are made, such as believing that a fast Fourier transform algorithm computes true Fourier transforms, or that any video features can be used for, e.g., segmentation, as long as the right classifier is used. Or, that edge or motion detection should always be the beginning of video analysis, old ideas that somehow have not completely faded.

I have found that these out-of-field students are generally marvelously hardworking, creative, and willing make the extra effort to climb a steep learning curve. The question I have asked myself is, how to adapt the course to serve such students, without compromising the rigor and depth of the topic? The answer that I have found works well is to offer a shallow curve at the outset (really, the first half of the course), proceeding through such introductory material as histograms, binary morphology, and the basic Fourier transform with great care, then, when the student is comfortable, accelerating the course and steepening the curve significantly, eventually covering advanced topics such as anisotropic diffusion, computational stereo, and video compression. I have been greatly satisfied by the results, and feel that the course has been significantly enriched by the astronomers, petroleum engineers, botanists, biomedical engineers, computer scientists, neuroscientists, and many other types who have taken it. I have also been impressed by the results. For example, one student from the Department of Kinesiology and Health Education at the University of Texas (UT) at Austin did a marvelous class project (I will talk about projects later) on spinal imaging, won the Best Project Award (voted on by the students, mostly engineers), and went on to publish her work [1].

Naturally, every teacher of video processing will want to tailor their course differently; some will choose to emphasize a deeper level of theory and a more rigorous treatment than I take, at least in the early parts of the course; however, I feel that the diversity I've discussed would be lost. So, here's my recommendation.

RECOMMENDATION 1

Make your video processing course accessible to students across engineering, science, and beyond. Yes, this means toning down the math, at least at first.

THE MULTIDISCIPLINARY NATURE OF VIDEO PROCESSING

A significant challenge facing the video processing educator is the deeply cross-disciplinary underpinnings of the subject (some of them shown in Figure 2). Video signals, after all, are the product of sensing an interaction between radiation (of some type; there are many that are used to form videos) and the environment. For the case of optical images, this implies at least some knowledge of how this interaction occurs, and how videos are sensed and transduced. What I think must be avoided is the idea that videos are "just arrays of numbers" to be played with. Indeed, too much of video processing in the past can be characterized as applying *ad hoc* methods to these arrays of numbers, without

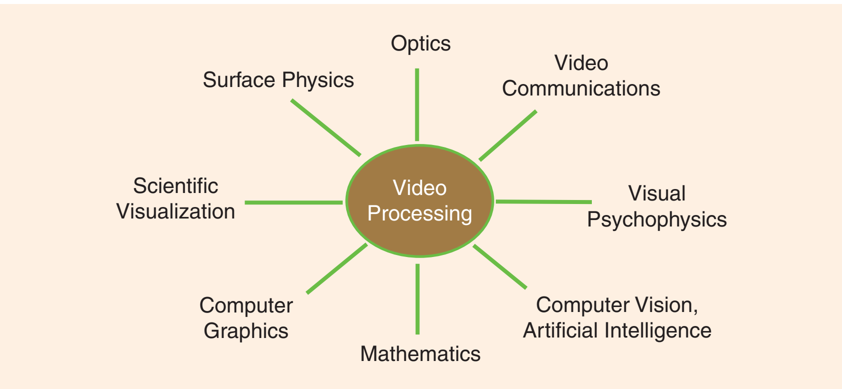

[FIG2] Some of the underpinnings of video processing.

making a nod towards the physics of video formation.

Of course, since covering all of the aspects of video formation constitutes a university course in itself, I find myself being selective in this regard. One thing that I try to avoid, however, is to merely mention the various aspects of video formation early on, then forget them—as many textbooks do. I think it's important that the properties of videos as they are formed survive into the processing steps. In other words, it's important that the video processing engineer remember and use information about the source of visual signals.

As an example of the kinds of principles that I try to cover in this manner, perhaps the most important is the simple perspective transformation—the geometry of optical image formation. This allows the student to understand the relationship between points in images and points in the real world. More important, I think, is conveying the massive loss of information that occurs by the process of projection from 3-D-to-two-dimensional (2-D) (or four-dimensional)(4-D)-to-3-D in the case of videos). As visual creatures endowed with brains able to extract enormous amounts of information from videos, students are often likely not to realize that much of the information they "see" resides in their brains, and not in the video itself—in the form of assumptions, internal models, world knowledge, and neural computations. I think it's important that the student understand that getting back the real-world information that is lost by perspective transformation is very hard! In the allied field of computer vision—the boundaries continue to blur between that field and this—this point is fundamental. Indeed, one might succinctly describe the field of computer vision as that of trying to reconstruct (and recognize) the real 3-D/4-D world from 2-D/3-D images/videos. I think it's important that this be given greater emphasis; I spend a lot of time on it early on in my own class, and return to it later, for example, in the context of computational stereo ranging.

Another good example is the loss of information arising from lens or other linear distortion. This leads to the topics of inverse filtering and restoration, essential but difficult aspects of video processing. While the student will understand that lenses blur, I choose to give real-world examples using data derived from an imperfect optical system, returning again to the principles of geometric optics to derive the blur function associated with that system. In this way, the student can be made to understand that there exist relationships between depth-of-field, aperture, and defocus.

Another basic principle that I think receives insufficient emphasis is the fact that video formation is a multiplicative interaction between radiant illumination and surface reflectivity. It is, perhaps, too much to ask to cover surface reflectance models in a basic video processing course; however, the student can be made to understand the multiplicative nature of video formation by a discussion of coherent sensing and multiplicative noise (e.g., speckle) [2]. The idea of a signal-to-noise ratio (SNR) that does not change with signal intensity is intriguing to the student, helps to establish this valuable facet of video formation in the students' minds, and in addition, leads naturally to the important ideas of logarithmic and homomorphic processing [3].

Another aspect of video processing that sets it apart is the great diversity of mathematics that is encountered. Aside from the traditional mathematics of signal processing, which are deep and diverse to start with, and include linear systems theory, complex analysis, power series, and transform and sampling theory, the field of video processing has been greatly enriched by theories involving partial differential equations, wavelets and multiscale, differential geometry, optimization theory, stochastic geometry, mathematical morphology, compressive sensing, and many more. To many students, this large and deep mathematical toolset can be daunting, but my approach to this is to explain early on that this is part of the richness and multidisciplinary nature of video processing, that many unsolved problems remain, and that these problems often attract some of the leading mathematicians in the world.

RECOMMENDATION 2

Emphasize the multidisciplinary nature of video processing to your student. After all, it is this nature that makes the topic so unique, interesting, challenging, and persistently vital.

THE RELEVANCE OF VISUAL PERCEPTION

It is on this topic that I might find myself in danger of being a pedagogue. Yet I feel that deepening student (and our own) understanding of biological visual perception is one of most important keys to creating future advances in video processing. We are, after all, visual creatures, and most videos that are "processed" automatically by computers are intended to be presented to human "receivers."

There are many ways in which I seek to engage the student's awareness and understanding of the perceptual aspects of video processing. Early on, I spend some time describing the human eye, including the retina and post-retinal cells, the nature of the fovea and the nonuniform sampling arrangement of the fovea, and the types and nature of eye movements. All of these are issues that come up later in the course. While many textbooks talk about the eye early on, very few carry these topics forward into discussion involving video processing, either from the perspective of optimizing processing for perception, or from the viewpoint that video processing can be made to emulate perceptual processes that have evolved, and ostensibly optimized, over the ages.

I think it is also extremely important that the student understand that the visual system is limited; not just in the usual senses of resolution in space or in time, but by the way processing occurs in the brain. With this in mind, I spend a significant amount of time (about one hour) talking about and presenting visual illusions. This has the added benefit of engendering a lot of interest, since most people are fascinated by the "magic" of visual illusions. But illusions are revealing as well, since they tell the student that "what you see" is not necessarily "what you get"—things in images and videos may be invisible, or transient in perception even though physically present and persistent. Conversely, things may be visible that do not exist, or that are different from what might be expected.

As vision scientists have learned much about visual processing in the last two decades, I think it appropriate to include visual models in video processing as much as possible. Indeed, many of the most significant advances in video processing and analysis have either come from vision science, or have been anticipated by what has been learned about visual functions.

Good examples of topics that connect the students' comprehension with biological vision include edge enhancement operators, such as the Laplacian-of-a-Gaussian operator, which closely model the receptive field profiles of post-retinal ganglion cells [4], Gabor functions, which are excellent image texture analyzers based on good models of the response characteristics of simple cells of visual cortex [5], the perceptual basis for Joint Photographic Experts Group/Moving Picture Experts Group quantization, spatial and temporal contrast sensitivity functions, stereopsis and random dot stereograms [6], and the perception of distortions (or assessment of quality) in images and videos [7], and sundry other topics, such as contour perception and optical flow. In all of these instances, using visual examples in the form of videos and images that have been processed, or that exemplify visual processes (such as the "random line" stereogram in Figure 3), is the most powerful way to convey the notions under discussion. Since video quality is a research specialty of mine, I usually involve the students in a human study as part of the course. This serves as a valuable research experience for them, while also supplying excellent subjects for our video quality studies!

As an advisor of graduate students (outside the classroom), I take this idea a step further by asking my students to attend classes in vision science, while also pointing them towards classic books [8] and broad, accessible articles [9] on various aspects of on visual perception.

RECOMMENDATION 3

Encourage your video processing students to become vision scientists.

MAKE IT VISUAL AND INTUITIVE

I would like to elaborate on the visual instructional aspect of teaching video processing. While video processing is founded upon deep perceptual and mathematical principles, explaining the math or the perceptual mechanism is usually not enough—the student needs to see the results of processing, or the perceptual effect. Literally, what they see is what they learn best. While scientific visualization adds value to the instruction of many sciences, in no field is this truer than in video processing.

Many of the popular texts on image and video processing recognize this by supplying visual examples throughout and offering software for the student to use [10]. However, in my view the key visual experience to be obtained is in the classroom, where dynamic visualizations can be provided using today's large-format digital displays and projectors. Further, I believe that it is necessary that this visual experience be highly interactive, so that examples of image and video processing can be done on a live basis, repeatedly, with the students seeing and participating in the selection of the processing parameters.

When I first began teaching this topic around 1990, I armed myself with hundreds of 35-mm slide examples to be projected as visual supplementary material to the algorithmic and theoretical topics being discussed. I felt that this provided some of the necessary intuition into the video processing that takes place and the consequent visual effects of that processing. As display and processing technology progressed, it became possible to develop interactive programs with easy-to-use graphical user interfaces (GUIs) that would allow me to project video processing algorithms and their results live on the large screen. This, I feel, was a vital development that enlivens and greatly deepens the intuitive visual experience that the student receives. It is far superior, I believe, to just plugging in visual examples (images and videos) to be displayed via Powerpoint or other presentation mode.

Towards this end, my students and I designed and created the Signal, Image and Video Audiovisualization (SIVA) gallery, which makes available an array of didactic tools for signal, image, and video processing education. We have explained the SIVA system, which is available at (http://live.ece.utexas.edu/class/siva/), in prior articles [11]; to date, the SIVA system is in use at more than 500 locations worldwide. Most important are the more than 100 demonstration programs written in MATLAB and in National Instruments Labview programs. These demonstration programs are especially designed for classroom instruction, as well as out-of-class use by students. Of interest here are the SIVA Image Processing and Video Processing galleries, which are written in Labview and are available free of charge by request.

The important elements of video processing demo programs for classroom use is that they have an easy-to-use and highly visual GUI, with clear controls for making parameter selections, and the ability to simultaneously display input and output images (as

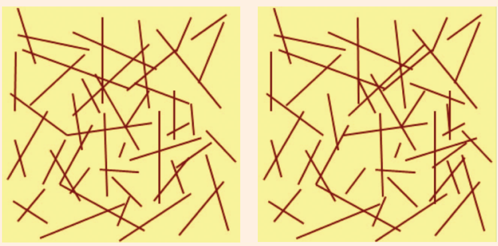

(a) (b)

[FIG3] Stereograms consisting of randomly placed simple elements such as (a) and (b) can help video processing students understand the science behind the hit movie *Avatar*.

applicable), with zoom for when the room is large. Naturally, the GUI and images should be easy to see (large enough images, fonts, and controls) when displayed on a liquid crystal display panel or by projection. Also, as I mentioned before, there should be the ability to rapidly show different results that are produced by varying the processing parameters, and also there should be an available library of several dozens of diverse images and videos containing different contents.

Figure 4 shows an example of the SIVA program interface for block truncation coding (BTC), a simple image codec that is still relevant for its simplicity and good performance, but in this context, more so for its explanatory power: the tradeoff between compression and image quality becomes clear as parameters that vary word length parameters are varied. The resulting compressed image and compression ratio are instantly displayed.

Figure 5 shows another SIVA demo program in action and exemplifies the theory of computing optical flow, which is an essential concept underlying most video processing algorithms. The algorithm implemented is the classic Horn–Schunck algorithm. The user is able to interactively control the number of iterations, the weight of the smoothing (regularization) term, and the values of parameters that dictate the way the resulting optical flow map is displayed as a needle diagram.

How does one develop visual teaching tools of this type? Naturally, one is free to utilize the SIVA system, which benefits from more than a decade of evolution and development. However, given one's own style of teaching, and one's own notions of what material should be covered, one may wish to develop video processing teaching tools of one's own design.

How to go about this? First, I urge that the courseware be highly visual, easy to use, interactive, and conducive to intuitive understanding of visual concepts.

Regarding the practical aspect of designing and creating courseware, I think that the best resource that you

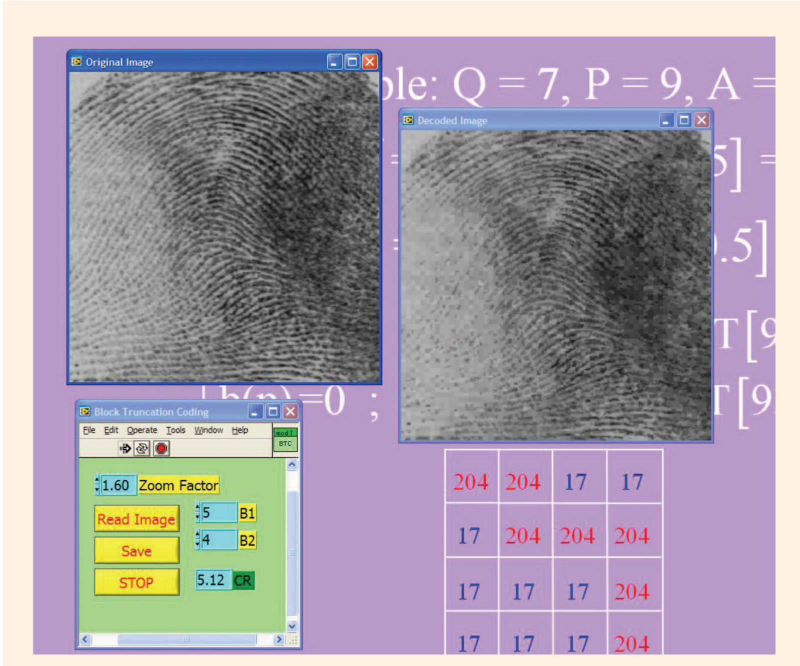

[FIG4] Example of SIVA program interface (overlaid over part of the classroom powerpoint from which it was linked) for BTC. Shown are original and compressed image, the choice of compression parameters, and the resulting compression ratio.

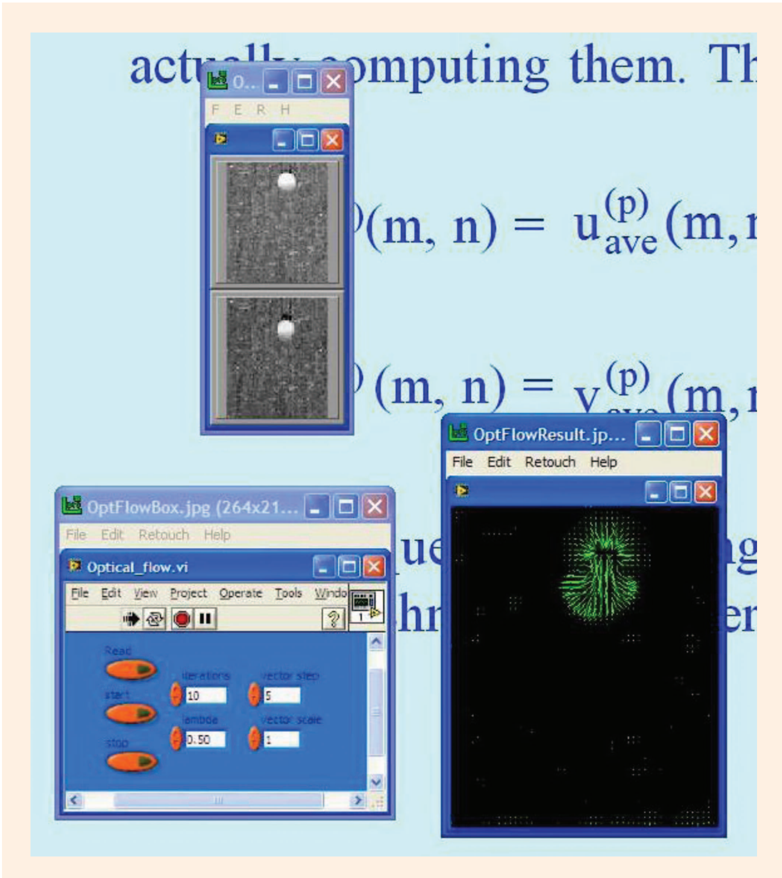

[FIG5] SIVA program interface (overlaid over classroom powerpoint) for optical flow computation. Also shown are two consecutive video frames and the computed flow.

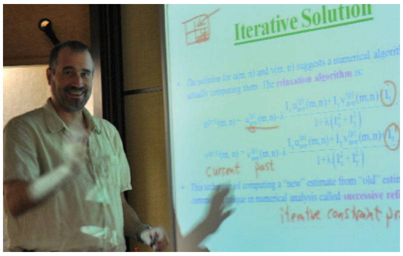

[FIG6] **The author teaching digital video processing to a present and remote audience using a Smart Board.**

have, aside from yourself, is students. I have been amazed and delighted by the energy and altruistic attitude that students take in helping to create courseware. For instance, I regularly receive e-mails from students that took my class many years prior, telling me how much they enjoyed the visual illusions part of the class, and that they had discovered a new and amazing illusion that I might want to include (and several times have). Indeed, the SIVA courseware would have been impossible to develop without heroic efforts from several students that I mention in the "Acknowledgments" section. Another significant resource is the makers of software tools that can be used to create courseware. National Instruments (NI), Inc., was unfailingly supportive in the development of the SIVA system in terms of both effort and resources.

RECOMMENDATION 4

Use highly visual and interactive teaching tools to teach the intuition behind video processing. If you want to create your own, get (lots) of help from students and software makers.

VIDEO HOMEWORK AND PROJECTS

Student participation in a video processing course by problem solving is quite important, of course. As far as regular assignments go, it is important that they be visually intuitive as well as challenging the student's analytic understanding of the material. I find that the best mixture of analytic, programming, and visual problem solving can be accomplished by having the student operate on visual signals themselves, through the modus of writing simulation software. MATLAB is ideal for this purpose. Using MATLAB, students are able to implement simple optical flow, edge detection, or video enhancement routines with relative ease. I find that about four homework sets per semester are about right, with each set asking the student to solve three to four problems.

However, by far the most important element of student participation, I have found, is the video processing class project. There are several reasons for this. First, video data is quite large, and it takes time to develop the facility to acquire, process, interpret and display videos. A project enables them to accomplish something meaningful with digital video. I encourage students to start early on their projects, and to be ambitious, within reason. Project proposals and progress reports are important ways to maintain a reality check on the projects, since students have a tendency to bite off more than they can chew. There is a sense among the uninitiated (probably propagated by Hollywood movies) that one can do anything with video processing, provided one's computer is powerful enough.

I am always amazed and delighted by the outcomes of the class projects. I have found that the energy and brilliance is always there, and only needs a little encouragement, in terms of ideas and assistance, but most of all motivation. Therefore, at the outset of the course, I tell the students that they will present their projects to the class at the end of the semester, and that the class will "grade" them based on the project's ambition, what the student(s) learned, and on their presentation. The combined scores are used to decide "winners" of the Best Project Awards (two groups are selected). These winners receive a small cash prize (US$200). More importantly, they receive an "A+" on their project—and here's the kicker—they get an "A" on the final exam, which they are no longer required to take! I realize that by doing this I am unabashedly fostering a rather intense rivalry among the students, but it has remained very friendly and the results have been amazing. The scoring by the students has always been fair and on-target.

The projects that have been presented are usually very practical, often are demonstrated "live" by the student with camera and laptop in hand, and are a lot more fun than the IEEE International Conference on Image Processing, held annually in the fall. Projects have included face recognizers, people trackers (using motorized cameras), drunkenness analyzers (by analyzing videos of iris behavior), vision-guided robotic cars, automatic cartoon-making from videos, and innumerable others. In the current semester, several students are creating video-related iPhone apps! It's so much fun (and impressive) that I now advertise the demonstrations to the general electrical engineering student and faculty body.

RECOMMENDATION 5

Make the video processing class project a centerpiece of your course, and have the students present their projects at the end. Find creative ways to motivate them!

THE FUTURE OF TEACHING VIDEO

The constant evolution of computer processing and display technology makes it possible to continuously upgrade our ability to communicate concepts in all branches of education, but especially in the highly visual field of video processing education. As I write this article, I have been lecturing all semester to a geographically distributed audience via distance learning technologies, which includes satellite video broadcast of the class to remote sites. Aside from audio and video of me, the course lecture notes (about 600 Powerpoint slides) and Labview-based demos are also broadcast and displayed on separate large-format screens. My current setup includes a Smart Technologies, Inc., Smart Board so that I can write directly on the slides and on the results of image/video processing, instantly visible everywhere the class is broadcast. Figure 6 shows the author using this integrated system. Of course the class is video recorded, with the digital recordings made immediately available on the class Web site for review.

Needless to say, I am delighted by these technology developments, not only because of their convenience and expansive educational potential, but because all of it, the cameras, displays, broadcast system, video courseware, and video recordings, are examples of byproducts of digital video processing research. It is very satisfying to lecture using the technology I am teaching.

Looking ahead, there are more exciting developments, not the least of which is 3-D. The movie *Avatar* has raised public awareness of the amazing experiences to be found in cinematic 3-D video. More importantly, 3-D technology is going to significantly penetrate the broader consumer market soon—3-D televisions are already commercially available, and glasses-free auto-stereoscopic displays will soon be good enough (and cheap enough) for the home audience as well. These displays will also be found on handheld devices. We aren't to the point of Princess Leia calling for "Obi-Wan Kenobi" via holo-projection, but we aren't far either.

These 3-D technologies will be available in the classroom as well. Before long, 3-D classroom displays will not be uncommon, and given the exposure and commercial drive in this direction, 3-D video instruction (meaning 3-D topics) and 3-D instruction techniques (meaning teaching in 3-D) are obvious developments to look forward to.

I hope that I have been able to express the enthusiasm and joy I find in teaching digital video processing. As I approach 30 years as a professor, there are many things that I do not look forward to every day, but one thing I always anticipate is lecturing on digital video.

RECOMMENDATION 6

Toss the chalk and whiteboard marker! Use modern video acquisition, communication and interactive display technology to teach digital video processing!

ACKNOWLEDGMENTS

The SIVA courseware was largely created by three students: Umesh Rajashekar, George Panayi, and Frank Baumgartner. Panayi wrote nearly all of the still image SIVA demos under my direction, while Baumgartner wrote the video SIVA demos. Rajashekar wrote all of the audio and one-dimensional demos which are written in MATLAB. He has also maintained the SIVA system (long after leaving UT at Austin), created the SIVA Web site, and coauthored our papers and reports on SIVA.

NI, Inc., located here in Austin, was unfailingly supportive in the development of the SIVA system. They were kind enough to fund the efforts of Panayi and Baumgartner (both NI engineers as well as UT students), and have taken an active role in maintaining SIVA through every Labview upgrade. More recently, several NI engineers (Nate Holmes, Matthew Slaughter, Carleton Heard, and Nathan McKimpson) assisted Rajashekar and me in creating a stand-alone (compiled) version of the SIVA image processing demos, complete with global user interface (from which all demos are easily accessed), for inclusion on a CD-ROM with my recent book [12]. NI manager Dinesh Nair assisted us in overseeing these efforts and in Chapter 2 of [12] helped us explain the system and the use of Labview.

REFERENCES

[1] D. S. Teyhen, T. W. Flynn, A. C. Bovik, and L. D. Abraham, "Digital fluoroscopic video assessment of sagittal plane lumbar spine flexion," *Spine,* vol. 13, no. 1, pp. E406–E413, Jan. 2005.

[2] M. Tur, K. C. Chin, and J. W. Goodman, "When is speckle noise multiplicative?," *Appl. Opt.*, vol. 21, no. 7, pp. 1157–1159, 1982.

[3] A. V. Oppenheim, R. W. Schafer, and T. G. Stockham, Jr., "Non-linear filtering of multiplied and convolved signals," *Proc. IEEE,* vol. 56, no. 8, pp. 1264–1291, Aug. 1968.

[4] D. Marr and E. Hildreth, "Theory of edge detection," *Proc. R. Soc. Lond. B.,* vol. 207, no. 1167, pp. 187–217, Feb. 1980.

[5] A. C. Bovik, M. Clark, and W. S. Geisler, "Multichannel texture analysis using localized spatial filters," *IEEE Trans. Pattern Anal. Mach. Intell.,* vol. 12, no. 1, pp. 55–73, Jan. 1990.

[6] B. Julesz, *Foundations of Cyclopean Perception*. Chicago, IL: Univ. of Chicago Press, 1971.

[7] K. Seshadrinathan and A. C. Bovik, "Motion-tuned spatio-temporal quality assessment of natural videos," *IEEE Trans. Image Processing,* vol. 19, no. 2, pp. 335–350, Feb. 2010.

[8] S. E. Palmer, *Vision Science: From Photons to Phenomenology*. Cambridge, MA: MIT Press, 1999.

[9] M. Carandini, J. B. Demb, V. Mante, D. J. Tolhurst, Y. Dan, B. A. Olshausen, J. L. Gallant, and N. C. Rust, "Do we know what the early visual system does?," *J. Neurosci.,* vol. 25, no. 46, pp. 10577–10597, Nov. 2005.

[10] R. C. Gonzalez, R. E. Woods, and S. L. Eddins, *Digital Image Processing Using Matlab*. Gatesmark, 2009.

[11] U. Rajashekar, G. Panayi, F. P. Baumgartner, and A. C. Bovik, "The SIVA demonstration gallery for signal, image, and video processing education," *IEEE Trans. Educ.,* vol. 45, no. 4, pp. 323–335, Nov. 2002.

[12] A. C. Bovik, *The Essential Guide to Image Processing*. New York: Elsevier Academic Press, 2009.

[SP]

[applications CORNER] continued from page 111

REFERENCES

[1] M. van Lieshout, *Markov Point Processes and Their Applications*. London: Imperial College Press, 2000.

[2] W. Ge and R. T. Collins, "Marked point processes for crowd counting," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009, pp. 2913–2920.

[3] N. Krahnstoever and P. Mendonca, "Bayesian autocalibration for surveillance," in *Proc. IEEE Int. Conf. Computer Vision,* Oct. 2005, pp. 1858–1865.

[4] D. Rother, K. A. Patwardhan, and G. Sapiro, "What can casual walkers tell us about a 3d scene?," in *Proc. IEEE Int. Conf. Computer Vision,* Oct. 2007, pp. 1–8.

[5] A. J. Baddeley and M. vanLieshout, "Stochastic geometry models in high-level vision," in *Statistics and Images,* vol. 1, K. V. Mardia and G. Kanji, Eds. Nashville, TN: Abingdon, 1993, pp. 231–256.

[6] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua, "Multi-camera people tracking with a probabilistic occupancy map," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. 30, no. 2, pp. 267–282, 2008.

[7] T. Zhao and R. Nevatia, "Bayesian human segmentation in crowded situations," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* June 2003, pp. 459–466.

[8] P. Green, "MCMC in image analysis," in *Markov Chain Monte Carlo in Practice,* W. R. Gilks, S. Richardson, and D. Spiegelhalter, Eds. London, U.K.: Chapman & Hall, 1995, pp. 381–400.

[9] H. Rue and M. Hurn, "Bayesian object identification," *Biometrika,* vol. 86, no. 3, pp. 649–660, 1999.

[10] X. D. M. Ortner and J. Zerubia, "A marked point process of rectangles and segments for automatic analysis of digital elevation models," *IEEE Trans. Pattern Anal. Machine Intell.,* vol. 30, no. 1, pp. 105–119, 2008.

[11] X. D. M. Ortner and J. Zerubia, "A marked point process of rectangles and segments for automatic nalysis of digital elevation models," *IEEE TPAMI,* vol. 30, no. 1, pp. 105–119, 2008.

[12] K. Otsuka and N. Mukawa, "Multiview occlusion analysis for tracking densely populated objects based on 2-D visual angles," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2004, pp. 90–97.

[SP]