Language Models for Everyone—Responsible and Transparent Development of Open Large Language Models

Alessio Cavicchi

***Abstract*—Affordable infrared (IR) cameras provide a unique opportunity to motivate and enhance the education of engineering students. We used both IR and visible images from a FLIR E60 camera as an instructional vehicle in a digital image processing course, where students came from several engineering majors. These IR and visible images were used as the basis for an open-ended final project in the course. Assessment via both pre- and post-project questionnaires showed the project was a positive experience for the students, and helped motivate them to learn the material. This paper discusses the course, the camera, the project, and how effective it was to add this project to the course.**

***Index Terms*—digital image processing, infrared, image fusion, engineering education**

I. Introduction

Digital signal processing (DSP), and the closely related area of digital image processing (DIP), have become very popular courses for many engineering students. Once mainly intended for electrical and computer engineering (ECE) students, the pervasive use of DSP and DIP in a wide range of research and industrial areas has caused other engineering and science majors to seek out these courses. If students have a solid mathematical background and a familiarity with topics such as linear systems theory, they can master the material in DSP and DIP courses, albeit sometimes with a higher degree of required effort than an ECE major. Having a class with a more diverse student background such as this increases the challenge for the professor to help the students learn the material.

For many research projects, there is a need to acquire data using cameras or other imaging systems, and usually also a need to manipulate and process this image data in various ways. This is no longer just an area for ECE students, and we frequently have students from most other engineering departments, along with students from some of the science departments (such as Molecular Biology, Physics, and Kinesiology), who need to master these concepts just to perform their research. To support this need, the ECE Department at the University of Wyoming offers four courses listed below, taught as a sequence by the same professor to ensure consistency across the courses.

1) Digital cameras, image formation, and image acquisition (graduate-level),
2) Digital image processing (senior-level and graduate-level combined section),
3) Advanced digital image processing (graduate-level), and
4) Object and pattern recognition (graduate-level).

This paper concentrates on course number 2 (the first of the two digital image processing courses), and in particular focuses on the final project assigned in the course. Note that a shorter, less detailed, description of this work, was described in [1].

In general, engineering educators tend to agree that interactive learning, exercises, and demonstrations are very useful for helping students to build an initial mental model of a given concept [2]–[6]. Using previously learned mental models to segue to new concepts is very effective variation on this technique [7], [8]. However, demonstrations can only help so much; actual hands-on exercises and projects have been shown to be one of the most reliable methods that lead to improved student learning [9]–[15]. There are even specific books and websites that support this approach [16], [17].

With relatively recent price drops for high-quality infrared (IR) cameras, we have started to introduce the IR camera as a tool to incorporate into various courses and outreach efforts, in an attempt to increase student enthusiasm and motivation to learn the course material. [18] The open-ended project we describe in this paper used images from an IR camera.

II. The Course

The first of the two digital image processing courses we offer is provided as a senior-level and a graduate-level combined section, as some seniors request this topic for one of their undergraduate electives. The graduate students in the course have more requirements for the course projects, but otherwise both groups of students are treated equally. The text we use is the excellent Gonzalez and Woods book [19], and the first course covers approximately the first half of the book (the second course then finishes the book). The topics covered by the first course are:

- Introduction to Digital Image Processing **(Chap. 1–2)**
  - What is it, why do we do it? Origins and examples of image processing (DIP)
  - Fundamental steps in DIP, components of a DIP system
  - Perception, image formation and optics, pixelization, quantization
  - Basic pixel descriptions, measures, and operators
- Image enhancement in the spatial domain **(Chap. 3)**
  - Gray level transformations and histograms
  - Using arithmetic and logic operations on images
  - Spatial filters: smoothing, sharpening, other

- Image enhancement in the frequency domain **(Chap. 4)**
  - The 1-D Fourier transform, the DFT and FFT, in time and in space
  - The 2-D Fourier transform in space and its inverse
  - Filtering in the frequency domain: smoothing, sharpening, other
- Image restoration **(Chap. 5, up to Sec. 5.10)**
  - Types of image degradation
  - Spatial domain and frequency domain noise reduction
  - Reducing other forms of degradation: inverse filtering, Wiener filtering, other
  - Introduction to image reconstruction from projections (computed tomography)
- Final Project: current industry problem

There are computer projects and reports associated with each main topic of the course, with the programs all being written in MATLAB®. One point pertinent to the final project is that this course deals primarily with grayscale images. Color models and color images are presented in Chapter 6, and are therefore part of the second course.

The 30 students enrolled in the course for Fall 2016 included 18 seniors, all EE or Computer Engineering majors, and 12 MS/PhD students (8 EE majors and 4 ME majors).

III. The Camera

The infrared (IR) region of the electromagnetic spectrum has longer wavelengths than the visible light which humans can see. The *visible spectrum* spans wavelengths of approximately 400–700 nm [20], [21]. One representation of the electromagnetic spectrum is shown in Figure 1, where the IR wavelengths can be seen on the left of the visible spectrum.

Humans can't see IR; any IR energy can only be directly sensed by humans as heat. Thus, an IR camera must assign a range of "false" colors that we are able to see to represent the IR energy. This false color method is often called *pseudocolor* in the literature [19]. Capturing images in the IR wavelengths challenges camera designers. For example, normal glass appropriate for visible wavelengths tends to reflect rather than pass IR wavelengths, and normal camera imaging sensors (CMOS or CCD) have IR-blocking filters to prevent saturation [22]–[24]. This latter point is required because the silicon-based sensor arrays respond strongly to what are called the near-IR wavelengths, as can be seen in Figure 2.

These and other challenges have been overcome by camera designers, and IR cameras are now available that combine reasonable price with acceptable image measurement quality.

The particular camera we used in this course was the model E60 by FLIR [25]. See Figure 3. The E60, like many IR cameras available today, actually provides two independent cameras: a primary IR camera, and a secondary visible wavelength camera. Each camera has its own optics and image sensor, with each optimized for the appropriate wavelengths. Note in Figure 3 that the image resolution for the two cameras is quite different, since high resolution IR sensors would be very expensive. This difference, plus the different field-of-view (FOV) and the different optical axis for image formation

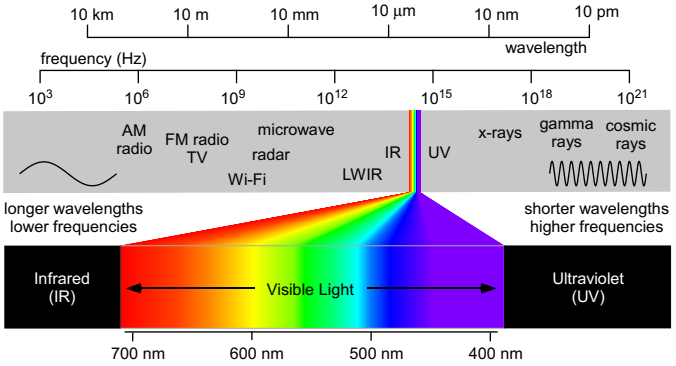

Fig. 1: Regions of the electromagnetic spectrum pertinent to this discussion.

between the two cameras, provide significant challenges for the final project in the image processing course.

A typical image from an infrared camera (a FLIR model C2) is shown in Figure 4, with the MSX image fusion option enabled. The model C2 is more compact than the E60, but also provides the MSX image fusion option. Note in Figure 4 that the key for the pseudocolor assignment of color to temperature (in degrees Celsius) is shown in a vertical bar on the right edge of the image. The temperature of the center of the image is displayed at the upper left of the image. Image fusion takes some salient features, such as edge information, from the associated visible image, and combines those features to provide an enhanced IR image. [26] Without image fusion, many IR images lack sufficient visual cues for the observer to know just what the image is showing.

IV. The Project

As an open-ended final project in our digital image processing course, we challenged the students to come up with their own version of a current engineering technique: how to use images from both an IR sensor and a visible-light sensor on the same unit to provide a greatly enhanced IR image using salient features (that they choose) from the visible-light image. This is a state-of-the-art technique, and is considered to be one type of "image fusion." [26] The proprietary method used by

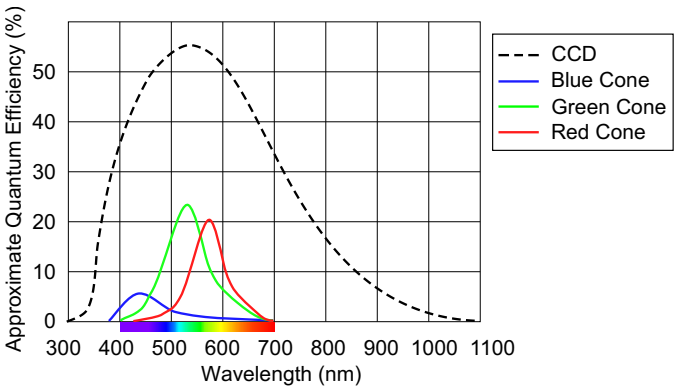

Fig. 2: Comparison of a typical silicon-based image sensor to human photoreceptors.

	FLIR E60
IR camera resolution	320 × 240
IR FOV	25°
IR spectral range	7.5 – 13 μm
IR temperature range	–20° to 650° C
Visible camera resolution	2048 × 1536
Focus	Manual
Image Fusion	Yes, MSX

Fig. 3: The model E60 IR camera from FLIR. Left image courtesy of FLIR, Inc. [25] Note: MSX is an image fusion technique proprietary to FLIR.

Fig. 4: A typical IR image taken with a FLIR C2, with the MSX image fusion option enabled. The image shows a kitchen backsplash on a wall with a poorly insulated fireplace behind it. The dark area in the upper left indicates a colder region, due to insufficient or missing insulation.

FLIR is called MSX, but the specific details of MSX were unavailable to the students. In order to make image fusion work for this scenario, an engineer must take into account the different resolutions, field of view (FOV), and other aspects that differ between the two types of image sensors. Coming up with a workable approach to image fusion is a difficult problem, and it draws upon all of the foundational image processing techniques the students have learned in the course. The students were told about this end-of-semester challenge project at the beginning of the semester, and its direct tie to current industrial practice both excited and motivated the students.

A simple scene, shown in Figure 5, was used to provide the sample images for the students. The IR image is lower resolution, much more narrow FOV, and was obtained using different optics which had an optical axis that was offset from the optical axis of the visible-light camera. One additional step was needed before the students could be given the images for the project. Since this first DIP course deals only with grayscale images, the images shown in Figure 5 were converted to grayscale (using a standard NTSC weighting scheme to convert RGB to luminance), as shown in Figure 6. Note the temperature annotations overlaid by the FLIR software on the IR image in Figure 5 have been turned off in Figure 6. The IR image on the right in Figure 6 is rather featureless, and is therefore an excellent example of why image fusion can enhance many IR images.

The students, in teams of two, were told to come up with a solution on their own for image fusion. The project instructions stated that their solution should be generalized to work with any scene, and not be a solution that works only for this one given scene. They were shown IR images with and without image fusion from the FLIR E60, and were allowed to use the FLIR E60 during several class periods to familiarize themselves with details about the camera and the images it takes. They were then provided the images shown in Figure 6 and turned loose.

V. Results

A. Image fusion

Overall, the results turned in by the 15 two-person student teams were impressive. While we readily answered questions posed by the students, great care was taken not to explicitly or implicitly imply any sort of preferred solution method. The students were free to choose how to co-align (i.e., register) the two very different images, pick what features to extract from the visible image, and how to "add" some part of those features to the IR image. In general, a common approach was to crop and shift, interpolate to a common resolution, detect edges, use the edges for alignment, extract edge information to be added, add a weighted version of the edges (using techniques similar to adding visible watermarks in the spatial domain) to the IR image, and then use some final image enhancement steps on the now-fused IR image. Some representative student results are shown in Figure 7.

Most team results were similar to the first three images shown in Figure 7, which are reasonably good enhancements of the original IR image. The difference comes from the

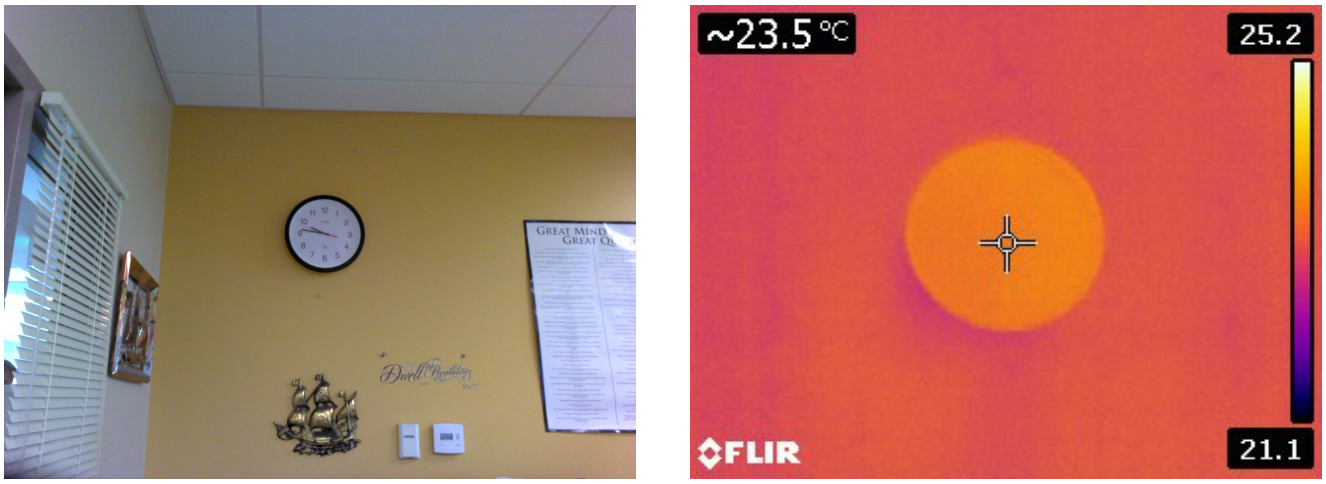

(a) Image from visible-light camera.

(b) Image from IR camera.

Fig. 5: Scene used to create sample images for the final project. Left: a 2048×1536 visible-wavelength image. Right: a 320×240 IR-wavelength image of the same scene.

(a) Image from visible-light camera.

(b) Image from IR camera.

Fig. 6: Grayscale versions of sample images for the final project.

different choices the students made. Image (d) is close to what the FLIR MSX technique would have produced. The worst result is shown on the bottom right of Figure 7. This team failed to get the images to line up (i.e., improper image registration) or even be the same FOV; you can see the IR version of the clock inside the superimposed clock detail. In their defense, this team claimed to have "run out of time" before they could perfect their method, but all teams had the same amount of time. The best result is shown on the bottom right of Figure 7, which is very close to what the proprietary FLIR MSX technique would have produced and is slightly sharper than image (d).

B. Assessment

To assess the effectiveness of this project for both helping the students learn and motivating them, two anonymous questionnaires were administered: one before the project, and one after the project. Both questionnaires had a 100% response rate of $N = 30$, and used a 5-point Likert scale shown below.

1) Strongly disagree
2) Disagree
3) Neutral
4) Agree
5) Strongly agree

The questionnaire before the project contained the following items, and Figure 8 shows the results.

1) I am already familiar with infrared cameras and images.
2) I already have a solid understanding of why image fusion is desirable for infrared images.
3) I already know how to implement image fusion for infrared images.
4) Being told we will be investigating the industrial challenge of image fusion for infrared images motivates me to better understand the topic of digital image

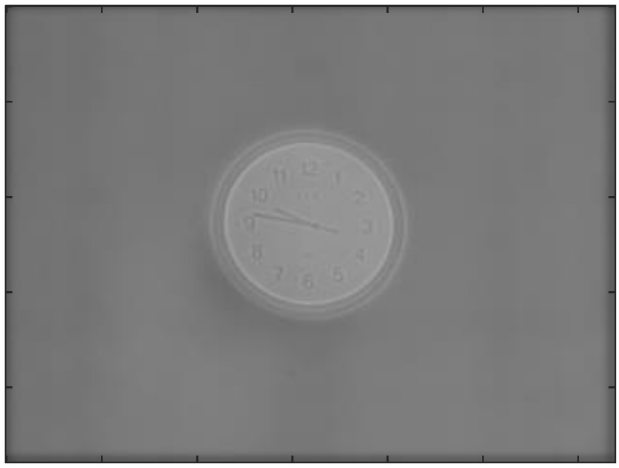

(a) A below average result: not enough edge emphasis.

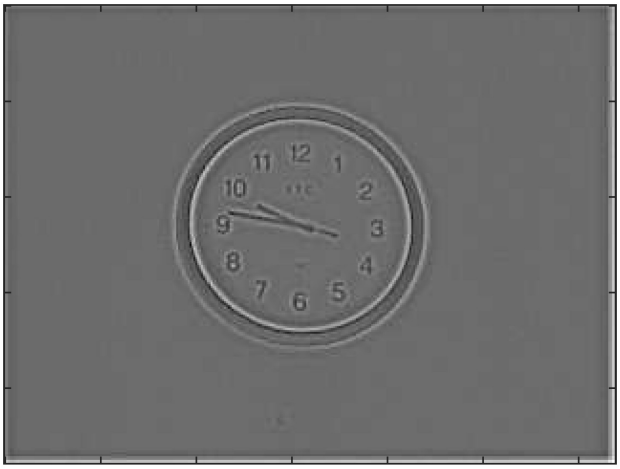

(b) A result with good edges but loss of contrast.

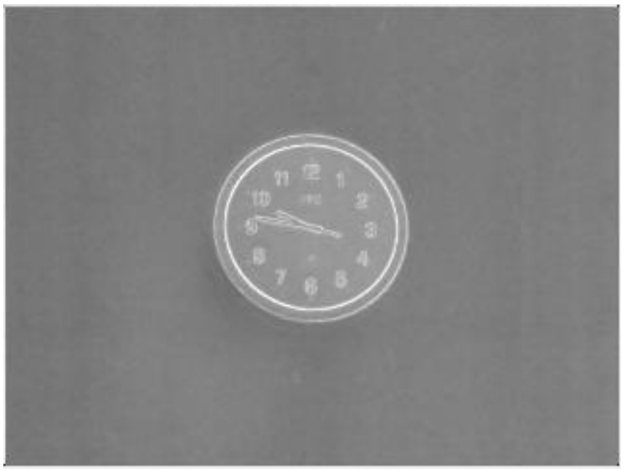

(c) A result with too much edge emphasis.

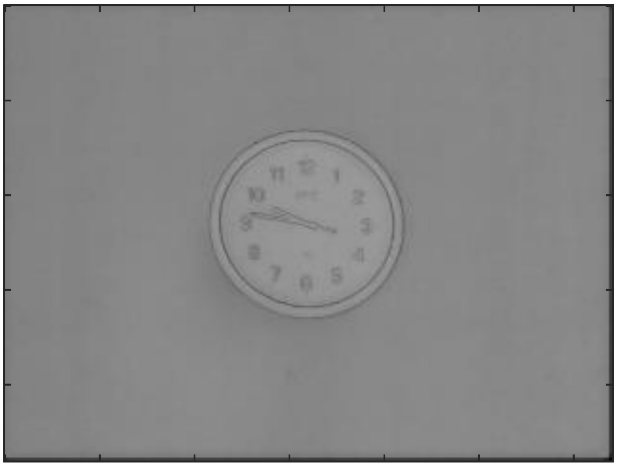

(d) A result close to what MSX would produce. One of the best.

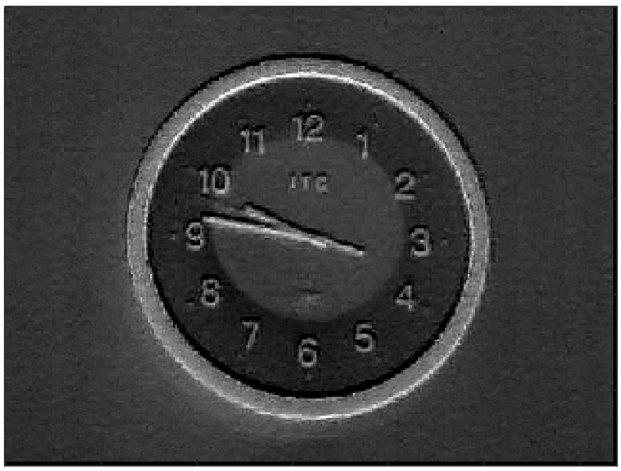

(e) The worst result; improper registration.

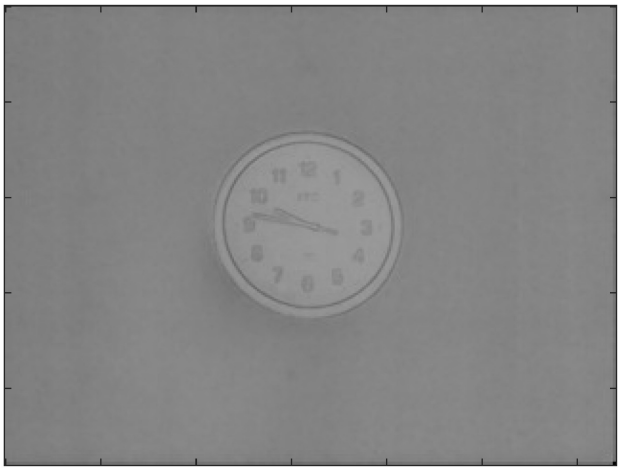

(f) The best result. slighty sharper than (d).

Fig. 7: Examples of student results for the final project. These are six representative results from a total of 15 teams that completed the project.

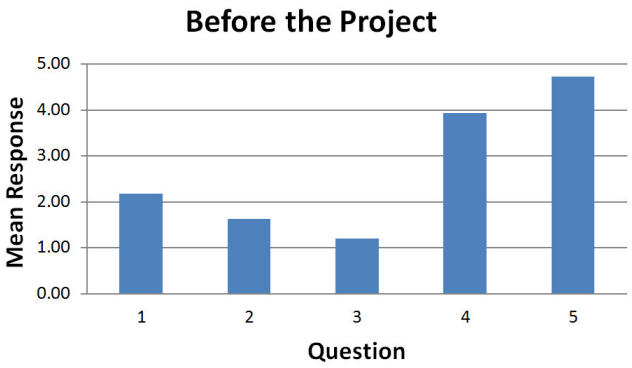

Fig. 8: The results of the "before the project" questionnaire.

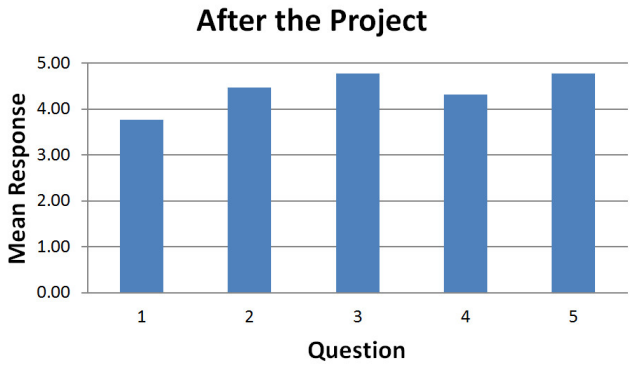

Fig. 9: The results of the "after the project" questionnaire.

processing.

5) Using current industrial challenges as the basis for a class project is a good idea.

The questionnaire after the project contained the following items, and Figure 9 shows the results.

1) I am now familiar with infrared cameras and images.
2) I now have a solid understanding of why image fusion is desirable for infrared images.
3) I now know at least one way to implement image fusion for infrared images.
4) Investigating the industrial challenge of image fusion for infrared images helped me better understand the overall topic of digital image processing.
5) Using current industrial challenges as the basis for a class project is a good idea.

As can be seen from the results, the questions that assessed how students learned new concepts as a result of the project (questions 1, 2, and 3) showed significant gain from before the project to after it. The questions that assessed the level of motivation the project provided (questions 4 and 5) were very high both before and after the project. Both of these results are what we had hoped to see.

VI. Conclusions

High quality, affordable IR cameras can be used in the classroom to provide an additional motivation for students to learn various concepts in engineering courses. We found that incorporating an open-ended final project, using IR images to solve a contemporary engineering challenge, provided enhanced learning and significant motivation for the students. This was confirmed by assessment data provided by questionnaires.

The course has been offered nine times since 2004, but this was the first time such an open-ended final project has been used. Anecdotally, the authors observed an obvious increase in excitement and enjoyment on the part of the students due to this project. We plan to continue to use such projects in the future.

References

[1] C. H. G. Wright and T. B. Welch, "Using IR cameras beyond outreach: motivational projects for engineering students," in *Proceedings of the 2017 ASEE Annual Conference*, June 2017. Paper AC2017-19308.

[2] C. S. Burrus, "Teaching filter design using MATLAB," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 20–30, Apr. 1993.

[3] R. F. Kubichek, "Using MATLAB in a speech and signal processing class," in *Proceedings of the 1994 ASEE Annual Conference*, pp. 1207–1210, June 1994.

[4] R. G. Jacquot, J. C. Hamann, J. W. Pierre, and R. F. Kubichek, "Teaching digital filter design using symbolic and numeric features of MATLAB," *ASEE Comput. Educ. J.*, pp. 8–11, January–March 1997.

[5] J. H. McClellan, C. S. Burrus, A. V. Oppenheim, T. W. Parks, R. W. Schafer, and S. W. Schuessler, *Computer-Based Exercises for Signal Processing Using* MATLAB 5. MATLAB Curriculum Series, Upper Saddle River, NJ (USA): Prentice Hall, 1998.

[6] J. W. Pierre, R. F. Kubichek, and J. C. Hamann, "Reinforcing the understanding of signal processing concepts using audio exercises," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 6, pp. 3577–3580, Mar. 1999.

[7] A. Rothenbuhler, C. H. G. Wright, T. B. Welch, and M. G. Morrow, "DSP see-through: Going beyond talk-through," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 2247–2251, May 2014.

[8] C. H. G. Wright, T. B. Welch, and M. G. Morrow, "Using student knowledge of linear systems theory to facilitate the learning of optical engineering," in *Proceedings of the 2015 ASEE Annual Conference*, pp. 26.1683.1–26.1683.11, June 2015. DOI: 10.18260/p.25019.

[9] C. H. G. Wright and T. B. Welch, "Teaching real-world DSP using MATLAB," *ASEE Comput. Educ. J.*, pp. 1–5, January–March 1999.

[10] T. B. Welch, M. G. Morrow, and C. H. G. Wright, "Teaching practical hands-on DSP with MATLAB and the C31 DSK," in *Proceedings of the 2000 ASEE Annual Conference*, June 2000. Paper 1320-03.

[11] C. H. G. Wright, T. B. Welch, D. M. Etter, and M. G. Morrow, "Teaching DSP: Bridging the gap from theory to real-time hardware," *ASEE Comput. Educ. J.*, pp. 14–26, July–September 2003.

[12] T. B. Welch, R. W. Ives, M. G. Morrow, and C. H. G. Wright, "Using DSP hardware to teach modem design and analysis techniques," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. III, pp. 769–772, Apr. 2003.

[13] C. H. G. Wright, M. G. Morrow, M. C. Allie, and T. B. Welch, "Using real-time DSP to enhance student retention and engineering outreach efforts," *ASEE Comput. Educ. J.*, pp. 64–73, October–December 2008.

[14] T. B. Welch, C. H. G. Wright, and M. G. Morrow, "The DSP of money," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 2309–2312, Apr. 2009.

[15] C. H. G. Wright, T. B. Welch, and M. G. Morrow, "Leveraging student knowledge of DSP for optical engineering," in *Proceedings of the 2015 IEEE Signal Processing and Signal Processing Education Workshop*, pp. 148–153, Aug. 2015.

[16] T. B. Welch, C. H. G. Wright, and M. G. Morrow, *Real-Time Digital Signal Processing: From MATLAB to C with C6x DSPs*. Boca Raton, FL (USA): CRC Press, 3rd ed., 2017.

[17] "RT-DSP website." http://www.rt-dsp.com.

[18] T. B. Welch, C. H. G. Wright, T. N. Kimmey, A. V. Delgado, S. O'Rorke, M. Brimstein, A. G. Norris, D. Buckmiller, R. Schwartz, D. R. Welch, and R. J. Edwards, "Seeing in the dark and through walls: using IR cameras in outreach," in *Proceedings of the 2016 ASEE Annual Conference*, June 2016. Paper AC2016-15631.

[19] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*. Upper Saddle River, NJ (USA): Prentice Hall, 3rd ed., 2008.
[20] R. Sekuler and R. Blake, *Perception*. New York: McGraw-Hill, 3rd ed., 1994.
[21] J. D. Bronzino, *The Biomedical Engineering Handbook*. Boca Raton, FL (USA): CRC Press, 1995.
[22] A. Daniels, *Field Guide to Infrared Systems, Detectors, and FPAs*. Bellingham WA (USA): SPIE Press, 2nd ed., 2010.
[23] M. Bass, ed., *Handbook of Optics*, vol. I. New York: McGraw-Hill, 2nd ed., 1995.
[24] G. C. Holst and T. S. Lomheim, *CMOS/CCD Sensors and Camera Systems*. Bellingham, WA (USA): SPIE Press, 2nd ed., 2011.
[25] "FLIR website, Exx-series IR cameras." http://flir.com/instruments/exx-series/.
[26] "Infrared resolution and contrast enhancement with fusion," 2013. U.S. patent 8,520,970 and 8,565,547.