

NOTE: this paper, titled

Towards commoditized smart-camera design

by

Boštjan Murovec, Janez Perš, Rok Mandeljc, Vildana Sulić Kenk and Stanislav Kovačič.

was published in

Journal of Systems Architecture, Volume 59, Issue 10 (part A), pages 847-858, 2013.

The version you have downloaded is in pre-print formatting. For the official version please see the publisher's web site:

<http://dx.doi.org/10.1016/j.sysarc.2013.05.010>

Towards commoditized smart-camera design

Boštjan Murovec*, Janez Perš, Rok Mandeljc, Vildana Sulić Kenk,
Stanislav Kovačič

*University of Ljubljana, Faculty of Electrical Engineering, Tržaška cesta 25, SI-1000
Ljubljana, Slovenia*

Abstract

We propose a set of design principles for a cost-effective embedded smart camera. Our aim is to alleviate the shortcomings of the existing designs, such as excessive reliance on battery power and wireless networking, over-emphasized focus on specific use cases, and use of specialized technologies. In our opinion, these shortcomings prevent widespread commercialization and adoption of embedded smart cameras, especially in the context of visual-sensor networks. The proposed principles lead to a distinctively different design, which relies on commoditized, standardized and widely-available components, tools and knowledge. As an example of using these principles in practice, we present a smart camera, which is inexpensive, easy to build and support, capable of high-speed communication and enables rapid transfer of computer-vision algorithms to the embedded world.

Keywords: commoditized smart camera, general-purpose smart camera,

*Corresponding author

Email addresses: `bostjan.murovec@fe.uni-lj.si` (Boštjan Murovec),
`janez.pers@fe.uni-lj.si` (Janez Perš), `rok.mandeljc@fe.uni-lj.si` (Rok Mandeljc),
`vildana.sulic@fe.uni-lj.si` (Vildana Sulić Kenk),
`stanislav.kovacic@fe.uni-lj.si` (Stanislav Kovačič)

1. Introduction

Capabilities of embedded processors have increased remarkably in recent years. Consequently, we have witnessed a migration of many tasks, previously considered as processing-intensive, to the domain of embedded systems. An example is computer vision, which is increasingly moving from the once-prevalent domain of desktop and industrial computers to the embedded devices — *embedded smart cameras* — and is the driving force behind the development of visual-sensor networks (VSNs) [1]. Especially in VSNs, *embedded vision* faces many challenges [2, 1, 3], due to severe limitations in computing performance, memory capacity and communication bandwidth of VSN nodes. Additionally, for VSNs to be a viable and economical alternative to centralized computer systems, nodes have to be inexpensive and easy-to-maintain. These requirements have resulted in numerous attempts to build an inexpensive yet sufficiently powerful *general-purpose* embedded smart camera, which would be able to run various computer vision algorithms.

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. Furthermore, no solutions are available at a price range that would justify large-scale VSN deployments with several dozens or even hundreds of nodes, collaborating on large, distributed and complex tasks. As shown by surveys of the field, e.g. [4, 3, 5], there is certainly no shortage of proposed smart camera architectures, which

24 all reached the stage of a working prototype. However, we feel that the ab-
25 sence of large-scale commercialization and deployment hints at more deeply
26 rooted weaknesses in existing general-purpose smart camera architectures.
27 We also feel that the absence of commercialization of *general-purpose smart*
28 *cameras* negatively affects the deployment of large-scale VSNs and therefore
29 warrants a special attention.

30 This paper aims to identify shortcomings of the existing designs. We
31 feel that the issues raised here significantly affect chances for widespread
32 commercial deployment of general-purpose smart cameras, especially in the
33 context of visual sensor networks. Therefore we also present a smart camera
34 design that, while by itself not being at the bleeding edge of technology,
35 does address many of the architectural weaknesses discussed in the paper.
36 As such, it should serve as an example of how the proposed principles may
37 be applied in practice, even when inexpensive hardware is used. We do
38 not claim that our proposition outperforms other designs; we merely claim
39 that we took care to address every aspect of the design in accordance with
40 the proposed principles, providing best possible environment for embedded
41 computer vision application design, *given the hardware constraints*.

42 The remainder of the paper is structured as follows: in Section 2 we
43 present a systematic overview of embedded camera designs and highlight
44 their strengths and weaknesses. Next, we describe the proposed set of de-
45 sign principles in Section 3, followed by the example of their application in
46 Section 4. In Section 5 we present a proof-of-concept camera design, which
47 follows our principles, along with the experimental evaluation. We conclude
48 the paper with Section 6.

49 2. Related work

50 A brief overview of embedded smart-camera technology is given in [6]; the
51 prevailing concept appears to be close integration of a CMOS visual sensor, a
52 microcontroller (MCU) and the supporting electronics [7], sometimes to the
53 point of integrating them all in a single integrated circuit, as illustrated in [8].
54 There is a wide variety of components that can be used in embedded smart
55 cameras, which result in different compromises between image resolution,
56 computing power, connectivity and power requirements. However, in this
57 paper, we focus on design principles behind a *general-purpose smart camera*.
58 Therefore, when comparing our work to state of the art, we attributed the
59 highest importance to the *programmer's view* of the camera. As such, we
60 divide the published camera designs into the following four categories.

- 61 • **Low-end.** Cameras from this group are, in general, built around ba-
62 sic, sometimes 8-bit, MCUs, and can process images at low framerate
63 and/or at low resolution. These cameras support only basic computer
64 vision algorithms, such as thresholding, simple blob extraction and
65 simple color segmentation. Consequently, their power requirements are
66 usually low and the hardware is inexpensive, but on the other hand,
67 they provide very little flexibility in terms of programming — the code
68 has to be written (and perhaps optimized) for each platform separately.
69 Often, such cameras lack even sufficient amount of memory to store the
70 whole image at once, which promptly disqualifies them for being ca-
71 pable of running anything but algorithms that require only single pass
72 over the image.

- 73 • **Mid-range.** Cameras from this group can employ a wide variety of
74 MCUs, but they do not run a fully fledged operating system. The two
75 main characteristics of cameras in this range are the availability of a
76 standard-compliant C compiler for the chosen MCU platform and suf-
77 ficient amount of RAM to store the whole image at a chosen resolution
78 and bit depth. Both characteristics enable a straightforward imple-
79 mentation of widely-used computer-vision algorithms for the camera
80 architecture. On the other hand, even though such designs are more
81 flexible, in the absence of an operating system, the task of such camera
82 remains more or less fixed once it is programmed.

- 83 • **High-end.** Cameras from this group run an operating system (fre-
84 quently a variant of Linux), offering high degree of flexibility. A stan-
85 dard C/C++ compiler is taken for granted, as is the number of stan-
86 dard software libraries like OpenCV [9] for computer vision, libjpeg
87 for JPEG image compression/decompression, and similar. Given the
88 proper connectivity, these cameras can be maintained and upgraded re-
89 motely, even to the point of completely changing their task (completely
90 different computer-vision algorithm) by simply instructing the camera
91 to run a different version of executable code.

- 92 • **Heterogenous and special designs.** Cameras from this group are
93 highly specialized for their task, and their hardware architecture reflects
94 that. As such, they cannot be considered *general-purpose*. From the
95 programmer's point of view, they may be more similar to low-end
96 designs regarding their inflexibility since the code has to be written from

97 scratch to take advantage of their capabilities, but their raw computing
98 performance at the task they have been developed for can be on a par
99 with or significantly better than the high-end designs.

100 *2.1. Low-end cameras*

101 There is a significant group of designs that are based on low-end MCUs.
102 In [10], authors present 30 frames-per-second color tracking, whereas in [11],
103 a tracking system with 7.5 frames-per-second is presented; both are based
104 on an 8-bit AVR MCU and use image resolution of 176×144 pixels. Both
105 designs have insufficient memory to store even a single full frame, and the
106 code was highly optimized to process image data in a single pass in order to
107 make even basic tracking feasible.

108 In [12], the problem was approached differently: a system-on-chip so-
109 lution was used, combining 8-bit AVR MCU with an FPGA module, and
110 external SRAM module is used to store acquired images. To take full advan-
111 tage of the architecture, processing has to be split between the FPGA and
112 MCU, requiring highly specialized and optimized code.

113 eCam [13] is an example of a low-end embedded camera design that is
114 specialized to the extreme; it consists only of image sensor, a specialized
115 controller chip (OV528 serial bridge with JPEG compression capabilities)
116 and communication interface (wireless sensor node). Its only function is
117 streaming of JPEG-compressed video to a base station.

118 *2.2. Mid-range cameras*

119 At the the lower end of this category, we find designs like a mote, described
120 in [14]. It uses 30×30 pixel optical mouse sensor and a dsPIC microcontroller

121 with only 16 kB of RAM, but is nevertheless able to perform Viola–Jones
122 face detection [15], gradient–based edge detection, motion estimation, and
123 background removal using two convex filters with slow and fast forgetting
124 parameters to avoid ghosting [16]. An ANSI C compiler is available for
125 the platform, and given the small resolution of image sensor, the amount of
126 memory is sufficient to store multiple acquired images. This way, computer–
127 vision algorithms can be adapted for use on this platform without excessive
128 optimization that is customary for low–end devices.

129 Cyclops sensor, presented in [17], is built around 8–bit Atmel ATmega128L
130 MCU, with RAM expanded to 64 kB and using the image sensor with max-
131 imum resolution of 352×288 pixels. In [18], a mote for wireless sensor net-
132 works is presented, built around an ARM7–based MCU with up to 64 kB
133 of RAM, and two 30×30 pixel image sensors. An ARM9–based camera is
134 presented in [19], processing 160×120 pixel images. With integrated energy
135 harvester (solar) module and a PIR sensor, the system is optimized for the
136 task of person detection, under the assumption of low–duty–cycle operation,
137 thus lowering power consumption. RoboVision platform in [20] comprises
138 ARM9–based MCU and 96 kB of RAM, but is coupled with disproportion-
139 ately high–resolution image sensor (1600×1200 pixels). Authors themselves
140 note that reasonable image sizes for onboard processing are actually much
141 smaller (160×120 pixels).

142 Finally, one of the most prominent embedded camera designs is the CMU-
143 Cam line, especially CMUcam3 [21]. Its ARM7–based MCU and 64 kB of
144 RAM are insufficient to process captured images at the maximum resolution
145 and bit depth offered by the included image sensor (352×288 pixels, RGB).

146 However, by halving the resolution and processing only grayscale images, it
147 can run state-of-the-art algorithms, such as Viola-Jones face detection [15].
148 A version of GCC-based toolchain, adapted to the camera architecture, is
149 provided.

150 *2.3. High-end cameras*

151 For computationally more demanding tasks, significantly more powerful
152 computing platforms are used, such as a RISC-processor-based design with
153 16 MB of RAM, presented in [22]. It is used for real-time car counting,
154 during which it processes images at the resolution of 320×240 pixels at 30
155 frames per second. Developers of the system explicitly aimed at a platform
156 without dedicated hardware. Even more powerful CITRIC platform [23] uses
157 Intel XScale PXA270 processor clocked at up to 624 MHz, combined with
158 64 MB of RAM and a 1280×1024 pixel image sensor; the platform runs
159 embedded Linux. Panoptes [24] relies heavily on standard, commoditized
160 hardware: Intel StrongArm 206 MHz processor (running Linux) is coupled
161 with a USB web camera (resolutions up to 640×480 pixels) and IEEE 802.11
162 (WiFi) wireless network interface.

163 *2.4. Heterogenous and special designs*

164 To cope with relatively high demands of real-time image processing, some
165 designs use specialized hardware to deliver required computing power. The
166 system presented in [25] relies on a DSP for real-time video compression, us-
167 ing optimized libraries provided by the DSP vendor. However, when running
168 plain C++ code, the system's performance suffers significantly.

169 There are many examples of similar systems: [26] and [27] rely on DSPs as
170 well, and [28] relies on specialized multimedia processors. Both [29] and [30]
171 rely on SIMD processors, requiring specialized programming knowledge to
172 exploit their full potential. On the other hand, the processing module is not
173 the only component that can be customized to optimize certain performance
174 goals: in [31], authors opt for a dual-sensor architecture to decrease power
175 consumption of an otherwise more powerful 32-bit microcontroller; they em-
176 ploy low-resolution imaging sensor for basic detection, and capture image of
177 higher resolution only when needed. Similar path is chosen by authors of the
178 dual-camera sensor, presented in [32].

179 It is obvious that the computer-vision algorithms need to be adapted
180 to take advantage of such multi-sensor setups, even if designs are based on
181 otherwise standardized architectures, and therefore are rightly considered
182 specialty designs.

183 *2.5. Summary*

184 Generally, high-end cameras employ more powerful hardware compo-
185 nents. Consequently, their cost and power requirements are higher as well,
186 but they are able to run more advanced computer-vision algorithms. How-
187 ever, there are some designs that stand out, for example [14], which has ex-
188 ceptionally modest hardware specifications, but is versatile and able to run
189 even state-of-the-art algorithms, albeit under obvious hardware constraints.
190 It is our opinion that such combination reflects a well-thought-out design
191 that provides high value to the user (programmer, software developer).

192 In this paper, we strive to formalize the design principles that would
193 have a similar effect: cost-effective, but maximally versatile general-purpose

194 embedded camera designs, which would be appealing to a wider user base and
195 not restricted to a single project or a single application. Finally, with regard
196 to the previously–described criteria, the camera that we present in this paper
197 as an practical example of applying the proposed design principles, falls into
198 the mid–range category.

199 **3. Proposed design principles**

200 We believe that there may be several design issues that hinder the wider
201 adoption of general–purpose smart cameras, especially in the context of visual
202 sensor networks. Those issues stem from certain design principles that may
203 be reasonable in general embedded system design, however, their justification
204 in the context of smart cameras is not entirely straightforward.

205 *3.1. Standardization and commoditization*

206 The reason behind the rapid development in many areas of computer
207 technology lies in widespread standardization, which allows unhindered com-
208 petition between many different vendors, and inevitably leads to commoditi-
209 zation of the new technology. A classical example is the personal computer,
210 which, after being designed and initially marketed by IBM, became extremely
211 commoditized, with its clones produced and sold by numerous vendors.

212 The second historically important aspect of such development is a possi-
213 ble inferiority of the winning product, which is again illustrated by the dom-
214 inance of the IBM–PC clones. At the time of its introduction, the PC was
215 an office computer with inferior graphics, no sound, and modest processing
216 capabilities. However, the effect of standardization was powerful enough to

217 overcome deficiencies and allowed rapid development into a far more superior
218 product.

219 At the moment, the field of smart cameras is experiencing a complete
220 absence of both paradigms, which contributes to slow adoption of embedded
221 visual systems and practically no major deployment of large-scale VSNs.
222 Such state of affairs is not surprising, as commoditization in short run benefits
223 the consumers of the technology and not its manufacturers. However, in the
224 long run, a widespread adoption of a particular technology often increases
225 manufacturer's production volume despite their lower market share.

226 *Our proposal.* Without the move to standardized and commoditized com-
227 ponents, it is unlikely that smart cameras and VSNs will ever become ubiqui-
228 tous. Furthermore, to spur the innovation in the field, especially the transfer
229 of state-of-the-art computer vision algorithms to the embedded domain,
230 and boost the popularity of smart cameras as much as possible, the designs
231 should be suitable for developers of various backgrounds and funding capac-
232 ities, from industrial teams to hobbyists.

233 In addition to being cost-effective, designs should extensively rely on
234 parts that can be easily purchased in small quantities, and allow assembly of
235 prototypes without highly specialized and expensive machinery. For exam-
236 ple, the components should be available both in SMD and DIP variation of
237 housing; the former is important for serial manufacturing, whereas the latter
238 is breadboard-friendly and enables quick and ad-hoc development.

239 As demonstrated by successful start-ups on a daily basis, the momentum
240 behind individual developers should not be underestimated. In our opinion,
241 such adoption is critical in evolution of smart cameras, as it significantly rises

242 the probability of *killer applications*, which are needed for the field of smart
243 cameras to stay competitive.

244 3.2. Long-term stability

245 The long-term fate of most designs is difficult to predict, since the rush
246 for specialized technologies usually results in use of the best and the hottest
247 parts that are available on the market at the instant of a project kickoff;
248 many such parts tend to be discontinued fairly frequently. As an example,
249 the renown line of CMUcam smart cameras¹ so far consists of four members,
250 with the first three already discontinued.

251 Frequent discontinuities and replacements pose a serious barrier to a wide
252 smart-camera adoption. Without considerable support from the camera de-
253 signers, computer-vision developers are forced to redesign their software in
254 order to cope with the changes introduced by new models. Such redesigns
255 are difficult to justify due to engineering costs for re-achieving something
256 that has already worked well, and there is always a risk of introducing re-
257 gressions. Despite the stated difficulties, a low-cost smart camera with a
258 long-term support is yet to appear.

259 *Our proposal.* There certainly exist smart camera applications that bene-
260 fit from the use of the latest technology; this is especially true in cases where
261 cameras perform relatively standard tasks, such as real-time video compres-
262 sion. Nevertheless, general-purpose smart cameras could definitely benefit
263 from the shift of focus from the bleeding-edge technology towards mature
264 and time-proven components, especially if they are to be deployed in tasks

¹<http://cmucam.org>

265 that will require many years of support and servicing.

266 3.3. Generality

267 The CMUcam family is also an illustrative example of products that
268 are non-generalized by design. Models CMUcam1² and CMUcam2³ were
269 specifically intended to be used as trackers, which was also reflected by their
270 firmware. Although it was possible to reprogram the whole camera to change
271 its capabilities, this was not their intended use, and such practice was not
272 officially supported. A further step in this direction has been done with
273 CMUcam4, where firmware is not even programmed in C or other widespread
274 programming language, but in a “C-like” language called SPIN⁴. Porting
275 existing computer-vision code to such environment is definitely not a trivial
276 task.

277 In contrast, the third member of the family, CMUcam3 [21], was fully-
278 programmable in standard C programming language; this model has become
279 extremely popular among computer-vision developers who required embed-
280 ded capability, which clearly demonstrates the need for general-purpose de-
281 signs. We expect that discontinuation of CMUcam3 will adversely affect
282 many research groups working with embedded vision.

283 *Our proposal.* The design of embedded smart cameras should be as
284 generic as possible. Instead of offering firmware with a pre-built function-
285 ality, the expectation should be that camera’s software will be developed
286 almost from scratch by *application* developers, not *camera* developers. The

²<http://cmucam.org/projects/cmucam1>

³<http://cmucam.org/projects/cmucam2>

⁴<http://www.cmucam.org/projects/cmucam4/wiki/Firmware>

287 latter should provide libraries and APIs for hardware-independent image ac-
288 quisition, image debugging and communication, preferably as C source code
289 that links with application's code, but should not impose any mandatory
290 boilerplate design of camera's firmware.

291 *3.4. Specialized technology*

292 Many designs rely on technologies that are too specialized, such as DSP
293 processors and FPGA circuits. Although these offer an unparalleled boost of
294 cameras' capabilities, very few computer-vision developers can actually use
295 them efficiently. This is directly tied to the available knowledge- and code-
296 base in the computer-vision community, which has at its access a large pile
297 of a carefully crafted C and Matlab code that uses only generic processing
298 hardware.

299 *Our proposal.* In contrast to the current practice, development of smart
300 cameras needs to be decoupled from development of applications that run
301 on them, since the two require different expertise. Computer-vision devel-
302 opment requires a dedicated computer-vision specialist. If such an engineer
303 needs to master additional specialty areas and hardware-specific skills, appli-
304 cation development becomes impractically demanding, since a whole-team
305 worth of expertise is required to develop even a simple practical solution. In
306 this context, truly open nature of an embedded camera becomes extremely
307 important. Cameras should be able to run standard computer-vision C code
308 nearly out of the box, and offer a simple C application programming inter-
309 face (API) for acquiring images directly into memory buffers specified by
310 developer. Different approaches almost inevitably confine camera's opera-
311 tion to one or at most a few applications that its own developers are willing

312 to implement.

313 3.5. Power supply

314 There appears to be consensus that embedded smart cameras need to
 315 be as energy-efficient as possible in order to enable battery-powered appli-
 316 cations, and the majority of designs nominally strive to achieve this goal.
 317 However, as shown in Table 1, most of the proposed solutions cannot oper-
 318 ate on battery power for long periods of time, especially when their power
 319 requirements in the fully-operational state is considered.

Design	Power [mW]	Autonomy [days]
[21] CMUcam3	650	0.57
[29] Xetal	600	0.6
[12] (WSN node)	500	1.33
[23] CITRIC	1000	1.5♣
[13] eCAM	231	1.6
[32] (dual-camera)	150◇	2.5
[30] IC3D	100	3.75
[18] (WSN mote)	99–363	1–3.8
[19] (with PIR sensor)	50–650◇	7.5–0.58
[17] Cyclops	23–65†	5.7–16†
[31] MeshEye	12‡	31‡

† rich power options to adapt to actual needs
 ‡ average; very low duty-cycle regime of operation
 ◇ low duty-cycle regime, wakeup by a secondary sensor
 ♣ authors' own estimation

Table 1: Power consumption of some of the designs when processing data (contrary to sleep mode) — exceptions are described in the footnotes. Where possible the consumption without communication was taken into account. Autonomy is based on 9 Wh capacity, which is an equivalent of two AA alkaline batteries without any additional energy source.

320 A general-purpose smart camera by definition cannot be designed with
 321 only one particular application in mind, therefore power calculations should

322 be based on its fully-operational state. Designs that conserve power by, for
323 example, resorting to extremely low frame rates or additional sensors for
324 wakeup, such as [31, 19, 32], thus allowing a camera to conserve power by
325 frequently shutting down major parts of its architecture, should be more
326 appropriately referred to as a low-duty-cycle *application* instead of a low-
327 power *design*.

328 Therefore, it seems that the ubiquitous goal of designing a general-
329 purpose battery-powered smart camera cannot be achieved using the cur-
330 rent technology. Even the extremely permissive scenarios with intermittent
331 operation (as envisioned in [31]) require monthly battery changes. In a large
332 network, this may be uneconomical in terms of battery and labor costs, un-
333 less there is absolutely no alternative. For a comparison, one should consider
334 the similar task of changing the burned-out light bulbs: the rush to substi-
335 tute incandescent light bulbs with alternatives was motivated in part by a
336 desire to reduce the labor costs, associated with frequent bulb changes. Fre-
337 quent battery changes may be impractical due to service interruption, and
338 environmentally harmful due to a pile of discarded batteries. Finally, even
339 though battery-powered designs seem attractive as a solution for a variety of
340 applications, there is rarely an absolute need to use battery power, at least
341 in urban environments – not to mention the indoor installations. In those,
342 the economics of the recurring cost of battery maintenance, compared to the
343 initially higher but fixed cost of wiring up the cameras, may quickly become
344 questionable.

345 Autonomy in the environments, where there is absolutely no possibility
346 of externally supplied power, can be extended using the solar power, such as

347 described in [19, 33], but this requires reserve capacity for the periods when
348 there is not enough solar radiation.

349 *Our proposal.* While there undoubtedly exist applications where battery-
350 powered operation is unavoidable, there certainly exist many more where
351 wired power works just as well. There are inherent advantages to the wired
352 power supply — mainly in relaxing other, often prohibitive, constraints, as
353 we demonstrate later on. Therefore, while the research into energy efficiency
354 of smart cameras remains a noble and worthwhile goal, there is no need to
355 restrict camera architectures solely to battery power, even in the case of
356 visual-sensor networks. To cope with broad range of situations, a truly open
357 and general-purpose camera architecture should support both power options
358 — the battery power and the wired power.

359 3.6. Communication

360 Another common design goal is wireless communication. In theory, battery-
361 powered and wirelessly-communicating nodes enable rapid deployment of
362 VSNs. In practice, the push for battery power necessitates resorting to
363 low-power communications, based on the IEEE 802.15.4 standard (ZigBee,
364 6LoWPAN), which have limited bandwidth and range. For example, in [30],
365 5 m is described as a practical operating distance; therefore, a square area
366 of 400 m² requires sixteen nodes just to overcome the limitations in com-
367 munication range, with associated sixteen battery packs to be changed on a
368 regular basis.

369 *Our proposal.* As with the power-supply issue, there certainly exist ap-
370 plications where wired communication is completely acceptable — even at
371 the price of a more complicated initial setup. In this case, the long-term

372 benefits of the wired solution are even more obvious, as wired communi-
373 cation enables significantly higher bit rates than low-power wireless solutions.
374 As we show later on, when a node is designed with wired power in mind, it is
375 trivial to extend its power-supply lines with additional high-speed commu-
376 nication lines at little extra cost. Additionally, wired communication is less
377 exposed to opportunistic hacking attacks as physical access is needed to tap
378 into wired communication lines.

379 When power lines are already in place but installation of communication
380 lines is unfeasible, the wired power supply may be efficiently supplemented
381 with a high-bandwidth wireless communication of a decent range (IEEE
382 802.11 WiFi). Of course, for nodes that *must* communicate wirelessly and
383 rely on battery power, low-power wireless communication is more or less the
384 only option.

385 When considering a truly open implementation of an embedded smart
386 camera, incompatible communication options may be difficult to overcome.
387 Many standard solutions, such as ZigBee, 6LoWPAN, WiFi and BlueTooth
388 exist in the form of modules that can be connected to standard SPI, I²C
389 or UART interfaces on MCUs; however, even in such cases, hardware from
390 different vendors likely requires different, vendor-specific, code paths in cam-
391 era's software.

392 In contrast, there are alternative solutions that do not require such ex-
393 plicit level of adaptation. This is especially true when standard UART com-
394 munication interfaces are used with less complex, time proven, communica-
395 tion options, such as RS-232 and RS-485, or a truly widespread and standard
396 interface, such as Ethernet. Consequently, communication tasks can be ab-

397 stracted away to the point where they do not represent an obstacle to a
398 widespread adoption anymore. Our proof-of-concept implementation shows
399 that this is indeed possible.

400 3.7. Image resolution

401 Existing embedded-camera designs frequently exhibit a noticeable dis-
402 crepancy between the resolution of imaging sensor and capabilities of the
403 associated MCU, and especially the amount of RAM available for image
404 storage and processing. Sensors often deliver color images of VGA or higher
405 resolutions, whereas storing a whole such image requires 1 MB of RAM or
406 more, which is well beyond the capacity of the low-end MCUs.

407 *Our proposal.* Many state-of-the-art computer-vision algorithms do not
408 rely on color information at all, and therefore a monochrome sensor is com-
409 pletely sufficient. Typical RAM capacities from 32 kB to 128 kB correspond
410 to monochrome images of resolutions from 50×50 to about 250×100 , which
411 should allow simultaneous storage of multiple images during processing. Ide-
412 ally, there should be support for image acquisition with adjustable resolution,
413 to enable online balancing between the available memory and desired visual
414 details.

415 Even though, intuitively, low resolutions seem insufficient for any com-
416 puter vision application, the contrary is successfully demonstrated by designs
417 like [14, 18] (30×30 -pixel applications). Another example is CMUCam1 with
418 the image resolution of 80×143 pixels [20]. Finally, if the resolution is too
419 high for the MCU that processes the data (to the point that the whole im-
420 age cannot be buffered in RAM), the camera is effectively downgraded from
421 the mid-range design to the low-end one, since the majority of ready-made

422 computer vision algorithms are prevented from running on such images.

423 *3.8. Image acquisition*

424 The main task of general-purpose smart cameras — and general-purpose
425 VSN nodes — is reliable image processing and extraction of information
426 required by downstream users (machines or people). Such concept is not
427 far from the field of machine vision, even if the setup does not operate in
428 an industrial environment. Consequently, the application of machine-vision
429 guidelines [34] may be beneficial for many potential applications.

430 One of the main guidelines in machine vision is the use of artificial illu-
431 mination for scene normalization, which is mostly avoided in smart-camera
432 designs due to power consumption of illuminators. The situation is contra-
433 dictory, since the lack of scene normalization requires more CPU-intensive
434 algorithms for image processing with adequately higher power consumption.
435 Another important lesson from the field of machine vision is that the optics
436 should be adapted to the problem at hand, thus making the best use of the
437 available image resolution.

438 *Our proposal.* We advise that the machine-vision guidelines are applied
439 whenever possible. The wired power, if feasible, is a game-changer in this
440 aspect; it on one hand enables use of illumination, which simplifies computa-
441 tion, and on the other hand allows use of a more powerful CPU than in the
442 case of a low-power design. Use of appropriate and possibly exchangeable
443 lens (in contrast to cameras with integrated lens, such as those usually built
444 into mobile phones) allows adaptation of optical-system parameters to a spe-
445 cific application. This includes, but is not limited to, compensating for the
446 previously-suggested lower image resolution by optimizing camera’s field of

447 view. Furthermore, optical filters may be used to increase camera's relative
448 sensitivity to a particular wavelength (typically wavelength of camera's own
449 illuminator), thus reducing the interference from uncontrolled light sources,
450 such as sunlight.

451 *3.9. Image debugging*

452 Development of computer-vision applications differs from the usual em-
453 bedded programming in at least one important aspect. During debugging,
454 the computer vision algorithm developer is not only interested in the plain
455 numerical values of MCU's registers and memory locations, but also benefits
456 from being able to display raw or partially-processed images stored in MCU's
457 memory. Many smart-camera designs place no emphasis on this functional-
458 ity, thus making embedded computer-vision development more challenging
459 than necessary.

460 *Our proposal.* Although image visualization is a higher-level concept than
461 inspection of memory locations, it is possible to make image debugging semi-
462 transparent by providing software interface for transferring image buffers
463 from the smart camera to the host PC, where they can be displayed. While
464 such solutions requires sufficient communication capabilities, they need to
465 be implemented only in the development version of the smart camera. We
466 demonstrate the concept of image debugging by piping live image buffers from
467 our proof-of-concept smart camera to the host PC and displaying them in
468 Matlab.

469 **4. Implementation of the proposed principles**

470 This section outlines possible implementations that reflect the previously
 471 discussed design principles. Following our classification of embedded-camera
 472 designs in Section 1, we focus on the low-end and mid-range category, as
 473 these represent significantly bigger challenge in that respect. Our recommen-
 474 dations are summarized in Table 2.

475 *4.1. Microcontroller and development toolchain*

476 The choice of MCU depends heavily on its integrated peripherals for
 477 CCIR image acquisition (Sections 4.4 and 4.5) and communication interfaces.

Aspect	Viable solutions
Microcontroller	Min. 64 kB RAM, 256 kB FLASH, two UART/serial modules. High-speed host USB module (only for UVC video). 32-bit CPU core with GCC-based or other standard C toolchain. C++ is desired but non-mandatory.
Power supply	A pair of exclusive switching-mode voltage regulators for wired and battery power (for each desired output voltage).
Communication	RS-485 for wired communication. Arbitrary wireless module with standard serial interface. Software abstraction for network independent API.
Imaging sensor	Analog CCIR (for monochrome low-resolution imaging) or USB camera supporting the UVC standard.
Video digitizer (for analog video)	Microcontroller's built-in A/D converter capable of at least 1 MSamples/s, together with other required periphery.
Optics	M12 interchangeable lens. Focal length depending on application.
Illumination	NIR LED diodes with optional diffusor and NIR filter.

Table 2: Standardized and commoditized technology for imaging and processing aspects of embedded smart-camera designs.

478 Multiple serial/UART interfaces are necessary to implement a variety of stan-
479 dard communication options, such as RS-232 and RS-485 (Section 4.3). The
480 majority of 6LoWPAN, ZigBee and Bluetooth modules and some Ethernet
481 interfaces can also be connected to serial interfaces.

482 MCU should have enough RAM to store all simultaneously-needed images
483 (Section 3.7) and other variables. In our view, 64 kB of RAM is a minimum,
484 but 128 kB is much better, especially if Matlab Coder⁵ is used to translate
485 Matlab code to the C language.

486 Since image processing is computationally intensive task, a MCU based
487 on a 32-bit CPU should be chosen; generally, it is less prudent to use 16-bit
488 or lesser CPU architectures, especially if power supply is not an issue. In
489 addition, MCUs that provide sufficiently advanced peripherals usually come
490 with 32-bit CPUs anyway, and wider buses also enable faster access to RAM
491 and peripherals. Hardware support for floating-point arithmetic is a bonus,
492 but not mandatory for many tasks.

493 CPU should be supported by a GCC⁶-based development toolchain to
494 facilitate straightforward porting of an existing computer vision code; both
495 native-C code, as well as code generated by Matlab Coder. Availability of
496 a C++ toolchain, which is less common in the embedded world, makes it
497 possible to use code based on OpenCV [9].

498 *4.2. Multiple power-supply options*

499 Notwithstanding the dilemma between the wired and battery power, it is
500 beneficial, if the camera can be powered from a variety of power sources. In

⁵<http://www.mathworks.com/products/matlab-coder/>

⁶<http://gcc.gnu.org>

501 addition to being flexible per se (allowing installation in variety of environ-
502 ments), such feature comes handy when autonomy is being extended using
503 rechargeable batteries, solar power (as for example in [33] and [19]) or pow-
504 ered from other unreliable sources, like car battery during engine cranking.

505 In the case of a wired power, the minimal upper end of input voltage range
506 should be 12 V, which is used by many illuminators. Furthermore, camera
507 operation on wider voltage ranges is easily achievable, which is needed for
508 seamless integration into various environments. For example, MAX5033⁷
509 voltage regulator supports input voltages between 7.5 V and 76 V; this in-
510 cludes 42 V, which is emerging as the new standard for car power supply.

511 *4.3. Communication solutions*

512 In contrast to frequently used low-power wireless solutions that mostly
513 benefit battery-powered nodes, there exist excellent wired alternatives. One
514 of them is the time-honored RS-485 standard, which is used in various in-
515 dustrial and consumer setups, but appears to be completely overlooked by
516 embedded smart-camera engineers. RS-485 has been devised for industrial
517 environments that require robustness to electromagnetic interference, dis-
518 tances up to 1.2 km and bandwidths up to 10 Mbit/s.

519 For example, members of MAX485 family⁸ offer bandwidth of 2.5 Mbit/s
520 and bus topology with up to 128 nodes at the price of \$3–5 in small quantities;
521 this is the price of a bare ZigBee transceiver integrated circuit, and one
522 third of a price of a full ZigBee transceiver module. Several members of

⁷<http://datasheets.maxim-ic.com/en/ds/MAX5033.pdf>

⁸<http://datasheets.maxim-ic.com/en/ds/MAX1487-MAX491.pdf>

523 MAX308x family⁹ provide bandwidth of 10 Mbit/s and bus topology with
524 up to 256 nodes for nearly the same price¹⁰. Both models exist in DIP and
525 SMD variants. Transceivers attach to standard UART interfaces. Connection
526 between RS-485 devices consists of a three-wire bus, where one of the wires
527 is a ground wire, and may be shared with power supply lines. In total, four
528 wires are enough to provide both power supply and RS-485 communication
529 to the node.

530 In our tests, we easily achieved bit rate of 3.5 Mbit/s over 125 m of
531 unshielded three-wire mains cable using MAX-485 transceiver that is de-
532 clared for maximal throughput of 2.5.Mbit/s¹¹. Although we certainly do
533 not suggest to use the transceiver beyond its specifications, the experiment
534 demonstrates that industry-certified and expensive RS-485 cables are not
535 necessary for embedded-camera setups, which makes RS-485 an extremely
536 attractive low-cost, high-range communication solution.

537 *4.4. Imaging sensor*

538 Due to popularity of digital-cameras and smartphones, many digital
539 imaging sensors are commercially available; however, each has a different,
540 non-standard, communication protocol. Such sensors are readily connectable
541 to MCUs, but if a sensor is withdrawn from the market, the embedded design
542 becomes obsolete. In addition, only high-volume customers can easily obtain

⁹<http://datasheets.maxim-ic.com/en/ds/MAX3080-MAX3089.pdf>

¹⁰The choice of bandwidth is more influenced by slew-rate limitation, EMI emission and length of connection than price.

¹¹Transmission-reception of 3 GB of data without a single bit error. Configuration consisted of one transmitter and one receiver. Practically achievable bandwidth degrades by increasing number of nodes but experiment nevertheless demonstrates a huge safety margin by exceeding the official bandwidth limits by 40%.

543 such sensors, whereas smaller groups often find such purchases challenging.
544 To the best of our knowledge, there are only two widespread standards for
545 imaging sensors that are not subject to the stated concerns: the time-honored
546 analog video¹², and the USB Video Class (UVC) specification¹³.

547 UVC cameras deliver images in digital format, and may thus seem more
548 suitable for embedded designs. However, although not required by the UVC
549 standard, most of available UVC cameras mandate high-speed (480 Mbit/s)
550 USB hosts, whereas many viable MCUs with built-in USB support can only
551 cope with full-speed (12 Mbit/s) USB connectivity. In addition, UVC cam-
552 eras generally deliver color images of high resolutions that require too much
553 RAM for storage and processing (see Section 3.7). Consequently, we see UVC
554 cameras as viable imaging sensors only in combination with high-end MCUs
555 that come with sufficient amount of RAM and high-speed USB support.

556 Considering all this, analog cameras present a viable imaging option. A
557 monochrome (CCIR) version is sufficient for many practical purposes (Sec-
558 tion 3.7). CCIR cameras of a size of a coin cost around \$6, and can be
559 obtained both in DIP or SMD housing.

560 4.5. Video digitizer

561 CCIR image of a sufficient quality for many computer-vision applications
562 can be acquired solely using MCU's internal peripherals, without any active
563 external circuits, such as analog filters and amplifiers. The required peripher-
564 als are an A/D converter and a voltage comparator with a voltage reference
565 for detection of sync pulses. An A/D with sampling rate of 1 Msample/s

¹²<http://pdfserv.maxim-ic.com/en/an/AN734.pdf>

¹³http://www.usb.org/developers/devclass_docs/USB_Video_Class_1.1.zip

566 allows acquisition at horizontal resolution of 50 pixels, whereas with faster
567 sampling it is possible to extend resolution up to the sensor’s physical limit.

568 Two adjacent interlaced video frames may be combined to double the
569 horizontal resolution at the same sampling rate, although the method is
570 suitable only for quasi-static images due to characteristic blurring that occurs
571 with fast moving objects. The full CCIR vertical resolution of 288 pixels
572 is achievable regardless of the A/D sampling rate. The stated capabilities
573 match the guidelines of Section 3.7.

574 Figure 1 shows two images acquired using the presented approach; resolu-
575 tions are 50×50 and 100×250 — the lower and the upper end of resolutions
576 recommended in Section 3.7.




Figure 1: Acquired images of a human hand without use of illuminators (left: 50×50 pixels, right: 100×250 pixels).

577 4.6. Optics, illumination and filters

578 According to the guidelines in Section 3.8, we recommend equipping cam-
579 eras with exchangeable M12-type lens, which are a standard for low-cost
580 (board) cameras. Several models cost around \$5 and fit the previously-

581 mentioned low-cost CCIR cameras. When the system is subjected to un-
582 wanted visible light, NIR LED illuminators in combination with visible-light
583 blocking NIR filter aid in scene normalization. Figure 2 demonstrates ben-
584 efits of this approach, which makes segmentation of an object of interest (a
585 hand) much easier.




Figure 2: Scene normalization with illumination. Separation of object of interest (a hand) from background becomes much easier compared to images in Figure 1, which were acquired using the same optics and at same resolution (left: 50×50 pixels, right: 100×250 pixels).

586 5. Proof-of-concept embedded smart camera

587 We illustrate the presented design principles on a proof-of-concept low-
588 cost embedded smart camera, which can be used as a standalone entity or in
589 role of a VSN node. Figure 3 presents the conceptual scheme.

590 5.1. Hardware

591 The MCU of our choice (U1 in Figure 3) is a high-end member of Mi-
592 crochip's PIC32 family¹⁴. It comprises 32-bit MIPS CPU, 512 kB of FLASH,

¹⁴<http://www.microchip.com/pagehandler/en-us/family/32bit/>

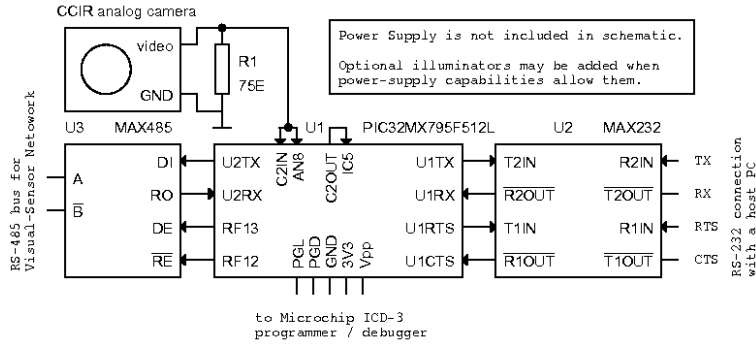


Figure 3: A proof-of-concept embedded smart camera.

593 128 kB of RAM, 10-bit A/D converter with 1 Msample/s, two voltage com-
 594 parators, a settable voltage reference, six UART modules, four SPI modules
 595 and five I²C modules. Microchip offers a free graphical integrated develop-
 596 ment environment MPLAB-X¹⁵ with assembler and C toolchain. The chosen
 597 MCU is available only in SMD housing, but it can still be used on breadboard
 598 with an aid of Sparkfun adapter¹⁶.

599 The preference for Microchip over other families and vendors is influenced
 600 by long-term stability of their products, especially in contrast to ARM pro-
 601 cessors, where different models are introduced and withdrawn fairly quickly.
 602 Also, the zero-price entry barrier for almost fully-featured development en-
 603 vironment is an important factor of consideration.

604 From a hardware standpoint, a plain MCU can be converted into a smart
 605 camera simply by adding a CCIR camera and resistor R1 (Figure 3), therefore
 606 enabling the capture of up to 50×250 pixel images at 25 frames per second,
 607 and up to 100×250 pixel images by combining two interlaced video frames.

¹⁵<http://www.microchip.com/pagehandler/en-us/family/mplab/>

¹⁶<http://www.sparkfun.com/products/9713>

608 The other two integrated circuits in Figure 3 add connectivity options to
609 the mote; part U2 takes care of RS-232 connection with the host PC¹⁷ for
610 image debugging and exchange of processing results, while part U3 drives
611 RS-485 bus network for VSN connectivity¹⁸. When wired power is used, it is
612 possible to use the RS-485 ground wire as power ground, therefore four wires
613 are enough to provide both power supply and connectivity to visual-sensor
614 network.

615 Our design requires extremely little specialized skills and can be built even
616 by hobbyists, which lowers the bar for low-volume applications. Total price
617 for the parts in Figure 3 together with omitted voltage regulators, lens and
618 illuminators, but without the PCB and housing, is about \$40 for purchases
619 in small quantities.

620 5.2. Software

621 Applications are developed in the standard C programming language.
622 Support for standard C code enables reuse of large code-base that exists
623 in the computer-vision community. The camera itself does not force any
624 boilerplate code or application skeleton for proper operation. In addition,
625 Matlab code can be straightforwardly ported to the camera by compiling into
626 C code using Matlab Coder. Our tests indicate that Microchip's toolchain
627 successfully compiles the resulting ports of Matlab code.

628 For decoupling application and camera development, we developed a soft-

¹⁷Depending on the actual RS-232 driver, a couple of additional discrete components not shown in the scheme may be needed.

¹⁸RS-485 bus must be terminated at both ends with 120 Ω resistor. When nodes' power supplies are floating against each other, an RS-485 ground wire must be added to the bus, and ground point of each node must be connected to the bus ground through 1 k Ω resistor.

629 ware library that offers platform- and imaging-sensor-independent API for
630 image acquisition. It is integrated into an application purely in form of source
631 code, without any binary objects. It offers C interface for image acquisition
632 into arbitrary RAM locations. The image resolution is specified at each ac-
633 quisition; it can be chosen from the resolutions that are available as per
634 Section 4.5. The library also handles acquisition of images with a prescribed
635 time delay, thus relieving the application developers from having to deal with
636 timers. Furthermore, it takes care of video signal sampling and operates as
637 a set of interrupt routines that run in the background, thus making CPU
638 available for execution of application's code even while the image acquisition
639 is in progress.

640 *5.3. Illustrative examples*

641 A simple test of image-processing capabilities was done using a motion
642 detector that acquires two 50×50 images with a prescribed time delay. For
643 each grabbed image, the following intermediate results are derived:

- 644 • image of vertical first-order derivatives (1.22 ms),
- 645 • image of horizontal first-order derivatives (1.21 ms),
- 646 • gradient image; per-element integer squared root of summed squares
647 of vertical and horizontal derivative (20.02 ms),
- 648 • blurred image; obtained from gradient image using 5×5 averaging filter,
649 together with scaling so that only additions are performed (2.58 ms).

650 The total execution time of the above steps is 25.03 ms; the timings are
651 for code compiled without optimization, which demonstrates the level of

652 performance offered by the free version of the toolchain.

653 The final image is produced by thresholding the absolute differences be-
654 tween two consecutive blurred images (1.60 ms). Motion is detected if the
655 sum of pixels in this image (0.59 ms) exceeds another prescribed threshold.
656 Please note that this is not a state-of-the-art motion detector but rather a
657 test-bed for timing common image-processing operations. According to the
658 results, the camera is capable of running the described detector at full 25
659 frames-per-second, and still have slightly more than 12 ms (30 %) of spare
660 time during each frame.

661 Computing variance of the resulting image (duration 31.94 ms) is also
662 an instructive measure of performance as it involves floating point arith-
663 metic. The chosen MCU does not offer hardware floating-point support,
664 which makes the computation rather slow. The last performance test in-
665 volving the motion detector is computing histogram of the resulting image
666 (duration 0.97 ms) and associated entropy using a look-up table together
667 with a proper scaling that enables operation in the integer domain (dura-
668 tion 0.13 ms). For illustration, the code optimization that comes with the
669 licensed version of the toolchain reduces image-processing steps from 25 ms
670 to 8.1 ms, whereas variance computation time drops from 32 ms to 17.7 ms.

671 The RAM capacity of the camera allows storage of all intermediate re-
672 sults, which makes it possible to transfer and examine them on the host PC
673 using the image-debugging facilities, as shown in Figure 4. Furthermore, we
674 have implemented two-way debugging facilities, which enable two images to
675 be uploaded from the host PC for processing on the camera. This way, algo-
676 rithms can be tested on a prescribed set of images (e.g., a standard dataset)

677 in the actual working environment instead on an emulator.




Figure 4: Image debugging of the motion detector. First row, from left to right: the first input image, vertical first-order derivatives, horizontal first order derivatives, gradient image. Second row: the second input image and its intermediate results. Third row: blurred variations in the first and second gradient image, thresholded difference of blurred images, histogram of the last image (with omitted frequency of pixels with zero value).

678 5.4. Additional tests

679 To illustrate the abilities of our proof-of-concept design, we ported sev-
680 eral well-known algorithms to the developed platform and measured their
681 performance in terms of processing time per frame. We also ran the same
682 tests on a range of hardware platforms to establish how they compare to our
683 smart camera.

684 *5.4.1. People tracker*

685 With only minimal modifications we successfully ported a basic people-
686 tracking application that was originally developed in the C programming lan-
687 guage under Linux. The algorithm is based on sequential image differencing
688 and blob tracking, with the Munkres assignment algorithm [35] as the final
689 tracking stage. The code on the camera runs at 20–25 frames-per-second at
690 resolution 50×50 .

691 *5.4.2. Object-recognition scheme and people detector*

692 We also implemented a simple object-recognition scheme, using either
693 HOG descriptor [36] or region covariance descriptor [37]. For each frame, it
694 performs a descriptor extraction and calculates its distance to the reference
695 descriptor. The recognition runs with 5.6 and 7.4 frames-per-second, re-
696 spectively (at resolution of 50×50 pixels). Finally, we run the combination
697 of the HOG descriptor and the linear SVM classifier, trained off-line as a
698 people detector (training was performed in advance on a PC). This test was
699 done at the resolution of 50×100 pixels – which is supported by the camera
700 as well. The test ran at approximately 2.5 frames-per-second. The detailed
701 timings are provided in the next section (Table 3).

702 Matlab code for covariance descriptor and the corresponding distance is
703 presented in Figure 5. It was ported to our platform using Matlab Coder.
704 Note that the code includes the computation of generalized eigenvalues, a
705 task that is seldom implemented on MCUs. We feel that the possibility
706 of running code that was developed in Matlab truly constitutes the Holy
707 Grail of the embedded computer vision, and illustrates the versatility of the
708 proposed approach.

```

% covariance descriptor for the
% central (32x32) region of a
% 50x50 pixel, 8-bit image.
function C = cov_descriptor (I)

% Copy, crop, convert
If = single(I(8:41,8:41));

% Convolution masks
f1 = [-1 0 1];
f2 = [-1 2 -1];

% Derivatives
Ix = conv2(f1, If);
Iy = conv2(f1,1,If);
Ixx = conv2(f2, If);
Iyy = conv2(f2,1,If);

% Features
If = If(2:33,2:33);
If = If(:);
coordinates = 1:32;
X = repmat(coordinates,32,1);
Y = repmat(coordinates',1,32);
X = X(:);
Y = Y(:);
Ix = Ix(2:33,3:34);
Ix = Ix(:);
Iy = Iy(3:34,2:33);
Iy = Iy(:);
Ixx = Ixx(2:33,3:34);
Ixx = Ixx(:);
Iyy = Iyy(3:34,2:33);
Iyy = Iyy(:);
F = [If, X, Y, Ix, Iy, Ixx, Iyy];

% Descriptor
C = cov(F);

% Distance between the two
% covariance descriptors
function D = distance (C1, C2)
D = sqrt(sum(log(eig(C1,C2)).^2));

```

Figure 5: Matlab code for computing the region covariance descriptor and generalized-eigenvalues distance. This code is directly portable to our camera using Matlab Coder.

709 *5.5. Comparison to widely-used hardware platforms*

710 In addition to testing our camera design, we ran the same battery of tests
711 on the following hardware platforms:

- 712 • **Axis 207W**, an ARM9TDMI-based IP camera running Linux (dis-
713 continued).
- 714 • **Axis P1346**, a HD IP camera with a CRIS ARTPEC-3 CPU. This
715 and the previous camera are commercial products, manufactured by
716 Axis Communications. They run Linux, and for licensing reasons, the
717 development toolchains are provided for both. We tested the cameras
718 without the video clients connected, therefore, their CPU load before
719 the test was negligible.
- 720 • **Raspberry Pi**, a single-board computer, intended as a tool for teach-
721 ing computer science, powered by ARM11-based ARM1176JZF-S pro-
722 cessor at 700 MHz, running Linux.
- 723 • **A high-end PC** with Intel Core i7 950 CPU at 3.07 GHz, running
724 Linux.

725 In all cases, GCC compiler for each platform was used to compile the same
726 C code, with full optimization enabled, except for the covariance test on Axis
727 207W, where the optimization produced broken code. In all cases, the code
728 utilized only a CPU; on PC, only a single core of a quad-core CPU was used.
729 Table 3 presents a comparison of results between all the tested architectures.
730 Note that we do not present the detailed results of the tracker test, as the
731 performance varies significantly with the number of objects detected. In all

732 cases, the camera draws 160 mA of current at 12 V without the infrared LED
 733 illuminator. This value includes the energy cost of voltage conversion. The
 734 whole setup consumes 360 mA with the illuminator turned on. (The mote
 735 without the illuminator can be powered from 5 V with an equal current
 736 consumption.)

Platform	Covariance [ms]	HOG [ms]	HOG+SVM [ms]
Axis 207W	2025 [†]	1845	3848
Axis P1346	262	261	574
Raspberry Pi	2.71	3.48	7.48
Intel PC	0.09	0.14	0.44
Our camera	135	180	395

[†] optimization disabled

Table 3: Comparison of the running times for the three algorithms on the five tested architectures. The values represent the average time needed for the actual computation. Covariance and HOG tests were run at 50×50 pixel resolution, while the people detector (HOG+SVM) was run at 50×100 pixels.

737 5.6. Exemplary communication abstraction

738 Section 3.6 mentions the idea of isolating core application’s code from
 739 communication details. To verify the feasibility of the approach, we devel-
 740 oped second smart-camera prototype by connecting a CCIR camera to Mi-
 741 crochip’s PIC32 Ethernet Starter Kit¹⁹, which consists of the same PIC32
 742 MCU combined with all hardware required for establishing an Ethernet
 743 connection. In conjunction with Microchip’s TCP/IP stack²⁰, such a node
 744 quickly becomes TCP/IP compliant.

¹⁹http://www.microchip.com/stellent/idcplg?IdcService=SS_GET_PAGE&nodeId=2615&dDocName=en545713

²⁰http://www.microchip.com/stellent/idcplg?IdcService=SS_GET_PAGE&nodeId=2505¶m=en535724

745 On both prototypes, a simple server application (written in C) services
746 requests from the host-PC client (written in Matlab), such as acquiring im-
747 ages at different resolutions, reading and writing of image buffers, sending
748 configuration information, etc. For both RS-232 and Ethernet, a set of rou-
749 tines for sending and receiving data, data availability testing, buffer-full state
750 indication, etc., was developed using the same function prototypes. For each
751 connection type, we prepared a C header file containing preprocessor macros
752 that translate into appropriate function calls.

753 The end result are two camera prototypes that, in spite of having two
754 radically different connection types, share the same application code. This
755 shows that isolation of embedded smart-camera applications from details of
756 communication options and hardware implementations is indeed both possi-
757 ble and beneficial.

758 **6. Conclusion**

759 In this paper, we identified several issues in the field of embedded smart
760 cameras, that in our opinion hinder widespread adoption of these devices.
761 Our research was partially motivated by our background in computer vision;
762 to us, a general-purpose embedded smart camera is a vehicle for unobtru-
763 sive and gradual introduction of computer-vision technology into everyday
764 use. Therefore, we are especially concerned with issues such as long-term
765 design stability, commoditization, generalization, rapid application develop-
766 ment and code reusability.

767 We certainly do not suggest that widely-discussed issues such as low
768 power and wireless communications do not require any attention from the

769 embedded camera community. However, we feel that there are significant
770 but overlooked opportunities for embedded camera designers to use proven
771 and well-established technologies to deploy computer vision and smart cam-
772 eras into widespread use. To illustrate our point, we presented our camera
773 design, which does not further state-of-the-art technology-wise, but on the
774 other hand possesses many properties we would like to see in a modern,
775 leading-edge smart camera. It is flexible in a sense that it allows the ap-
776 plication of machine vision principles to solutions based on such camera. It
777 allows reasonably fast communications, thus it can be used as a part of a
778 larger camera network. It enables computer-vision engineers to reuse a large
779 body of code developed for other platforms, and it allows computer-vision
780 scientists to quickly develop new algorithms using widely-used engineering
781 tools, such as Matlab. Finally, since it is made up of relatively standard or
782 widely used components, it should allow solution providers to deploy and sell
783 products based on such camera with relatively low risk of quick obsolescence.
784 We are not proposing our camera as a reference design for embedded smart
785 cameras, but we hope that the principles outlined in this paper gain wider
786 adoption in the embedded smart camera community and find their way into
787 next generation of cameras.

788 **References**

- 789 [1] S. Soro, W. Heinzelman, A Survey of Visual Sensor Networks, *Advances*
790 *in Multimedia 2009* (2009) 21 pages.
- 791 [2] Y. Charfi, N. Wakamiya, M. Murata, Challenging Issues in Visual Sensor
792 *Networks, IEEE Wireless Communications* 16 (2009) 44–49.

- 793 [3] I. T. Almalkawi, M. Guerrero Zapata, J. N. Al-Karaki, J. Morillo-Pozo,
794 Wireless multimedia sensor networks: Current trends and future direc-
795 tions, *Sensors* 10 (2010) 6662–6717.
- 796 [4] B. Tavli, K. Bicakci, R. Zilan, J. M. Barcelo-Ordinas, A survey of visual
797 sensor network platforms, *Multimedia Tools and Applications* (2011)
798 1–38.
- 799 [5] B. Rinner, T. Winkler, W. Schriebl, M. Quaritsch, W. Wolf, The evo-
800 lution from single to pervasive smart cameras, in: *Second ACM/IEEE*
801 *International Conference on Distributed Smart Cameras (ICDSC 2008)*,
802 pp. 1–10.
- 803 [6] Z. Zivkovic, R. Kleihorst, Smart Cameras for Wireless Camera Net-
804 works: Architecture Overview, in: H. Aghajan, A. Cavallaro (Eds.),
805 *Multi-Camera Networks*, Academic Press, 2009, pp. 497–510.
- 806 [7] B. Rinner, W. Wolf, An Introduction to Distributed Smart Cameras,
807 *Proceedings of the IEEE* 96 (2008) 1565–1575.
- 808 [8] W. Caarls, P. P. Jonker, H. Corporaal, SmartCam: Devices for Em-
809 bedded Intelligent Cameras, in: *PROGRESS 2002, Proceedings of 3rd*
810 *Seminar on Embedded Systems*.
- 811 [9] G. Bradski, The OpenCV Library, *Dr. Dobbs Journal of Software Tools*,
812 2000.
- 813 [10] H.-N. Nguyen, A.-C. Lee, Real Time Tracking Multiple YUV 24-Bit
814 Color Objects with 8-Bit MCU-Based Embedded Vision System, in:

- 815 Proceedings of the Fourth International Conference on Innovative Com-
816 puting, Information and Control (ICICIC'09), pp. 160–164.
- 817 [11] A. Aggarwal, Embedded Vision System, in: Proceedings of
818 IEEE/ASME International Conference on Mechtronic and Embedded
819 Systems and Applications (MESA'08), pp. 618–621.
- 820 [12] A. Kerhet, M. Magno, F. Leonardi, A. Boni, L. Benini, A low-power
821 wireless video sensor node for distributed object detection, Journal of
822 Real-Time Image Processing 2 (2007) 331 – 342.
- 823 [13] C. Park, P. H. Chou, ecam: ultra compact, high data-rate wireless sensor
824 node with a miniature camera, in: Proceedings of the 4th international
825 conference on Embedded networked sensor systems, SenSys '06, pp. 359–
826 360.
- 827 [14] M. Camilli, R. Kleihorst, Demo: Mouse Sensor Networks, the Smart
828 Camera, in: The Fifth ACM/IEEE International Conference on Dis-
829 tributed Smart Cameras (ICDSC'11), pp. 1–3.
- 830 [15] M. Viola, M. J. Jones, P. Viola, Fast Multi-View Face Detection, in:
831 Proceedings of Computer Vision and Pattern Recognition (CVPR'03).
- 832 [16] M. Cristani, M. Farenzena, D. Bloisi, V. Murino, Background Subtrac-
833 tion for Automated Multisensor Surveillance: A Comprehensive Review,
834 EURASIP Journal on Advances in Signal Processing 2010 (2010) 24
835 pages.
- 836 [17] M. Rahimi, R. Baer, O. I. Iroezi, J. C. Garcia, J. Warrior, D. Estrin,

- 837 M. Srivastava, Cyclops: In Situ Image Sensing and Interpretation in
838 Wireless Sensor Networks, in: SenSys, ACM Press, 2005, pp. 192–204.
- 839 [18] I. Downes, L. B. Rad, H. Aghajan, Development of a mote for wireless
840 image sensor networks, in: Proc. of COGNITIVE systems with Interactive
841 Sensors (COGIS).
- 842 [19] M. Magno, D. Brunelli, P. Zappi, L. Benini, A solar-powered video sensor
843 node for energy efficient multimodal surveillance, in: 11th EUROMI-
844 CRO Conference on Digital System Design Architectures, Methods and
845 Tools (DSD '08), pp. 512–519.
- 846 [20] V. Rana, M. Matteucci, D. Caltabiano, R. Sannino, A. Bonarini, Low
847 cost smartcams design, in: IEEE/ACM/IFIP Workshop on Embedded
848 Systems for Real-Time Multimedia (ESTImedia 2008), pp. 27–32.
- 849 [21] A. Rowe, A. Goode, D. Goel, I. Nourbakhsh, CMUcam3: An Open
850 Programmable Embedded Vision Sensor, Technical Report CMU-RI-
851 TR-07-13, Robotics Institute, Carnegie Mellon University, Pittsburgh,
852 Pennsylvania, 2007.
- 853 [22] M.-Y. Chiu, R. Depommier, T. Spindler, An Embedded Real-time Vi-
854 sion System for 24-hour Indoor/Outdoor Car-Counting Applications,
855 volume 3, IEEE Computer Society, 2004, pp. 338–341.
- 856 [23] P. Chen, P. Ahammad, C. Boyer, S.-I. Huang, L. Lin, E. Lobaton,
857 M. Meingast, S. Oh, S. Wang, P. Yan, A. Yang, C. Yeo, L.-C. Chang,
858 J. D. Tygar, S. Sastry, Citric: A low-bandwidth wireless camera net-

- 859 work platform, in: Second ACM/IEEE International Conference on
860 Distributed Smart Cameras (ICDSC 2008), pp. 1–10.
- 861 [24] W.-C. Feng, E. Kaiser, W.-C. Feng, M. Le Baillif, Panoptes: Scalable
862 low-power video sensor networking technologies, in: In MULTIMEDIA
863 03: Proceedings of the eleventh ACM international conference on Mul-
864 timedia, ACM Press, 2003, pp. 562–571.
- 865 [25] M. Bramberger, R. P. Pflugfelder, A. Maier, B. Rinner, B. Strobl,
866 H. Schwabach, A smart camera for traffic surveillance, in: In Pro-
867 ceedings of the First Workshop on Intelligent Solutions in Embedded
868 Systems, pp. 153–164.
- 869 [26] H. Andrian, K.-T. Song, Embedded CMOS Imaging System for Real-
870 Time Robotic Vision, in: IEEE/RSJ International Conference on Intel-
871 ligent Robots and Systems (IROS 2005), pp. 1096–1101.
- 872 [27] M. Bramberger, R. P. Pflugfelder, A. Maier, B. Rinner, B. Strobl,
873 H. Schwabach, A Smart Camera for Traffic Surveillance, in: Pro-
874 ceedings of the First Workshop on Intelligent Solutions in Embedded
875 Systems (WISES), pp. 153–164.
- 876 [28] T. Loten, R. Green, Embedded Computer Vision Framework on a Mul-
877 timedia Processor, in: Proceedings of 23rd International Conference on
878 Image and Vision Computing New Zealand (IVCNZ'08), pp. 1–5.
- 879 [29] R. Kleihorst, A. Abbo, B. Schueler, A. Danilin, Camera Mote with a
880 High-Performance Parallel Processor for Real-Time Frame-based Video

- 881 Processing, in: First ACM/IEEE International Conference on Dis-
882 tributed Smart Cameras (ICDSC'07), pp. 109–116.
- 883 [30] R. Kleihorst, B. Schueler, A. Danilin, M. Heijligers, Smart Camera Mote
884 with High Performance Vision System, in: ACM SenSys Workshop on
885 Distributed Smart Cameras.
- 886 [31] S. Hengstler, H. Aghajan, A Smart Camera Mote Architecture for Dis-
887 tributed Intelligent Surveillance, in: ACM SenSys Workshop on Dis-
888 tributed Smart Cameras.
- 889 [32] D. Xie, T. Yan, D. Ganesan, A. Hanson, Design and implementation of
890 a dual-camera wireless sensor network for object retrieval, in: Interna-
891 tional Conference on Information Processing in Sensor Networks (IPSN
892 '08), pp. 469–480.
- 893 [33] T. Celik, H. Kusetogullari, Solar-powered automated road surveillance
894 system for speed violation detection, *IEEE Transactions on Industrial
895 Electronics* 57 (2010) 3216–3227.
- 896 [34] B. G. Batchelor, P. F. Whelan, Machine Vision Systems: Proverbs, Prin-
897 ciples, Prejudices & Priorities, in: In Proceedings of Applied Machine
898 Vision Conference, pp. 7–19.
- 899 [35] H. W. Kuhn, The Hungarian Method for the Assignment Problem, in:
900 *Naval Research Logistics Quarterly*, volume 2(1–2), pp. 83–97.
- 901 [36] N. Dalal, B. Triggs, Histograms of Oriented Gradients for Human De-
902 tection, in: *IEEE Computer Society Conference on Computer Vision
903 and Pattern Recognition (CVPR'05)*, volume 1, pp. 886–893.

904 [37] O. Tuzel, F. Porikli, P. Meer, Region Covariance: A Fast Descriptor
905 for Detection and Classification, in: European Conference on Computer
906 Vision (ECCV), volume 7.



907 Boštjan Murovec received B.Sc., M.Sc. and Ph.D. degrees in Electrical engineering at
908 the Faculty of Electrical Engineering, University of Ljubljana, in 1996, 1999 and 2002,
909 respectively. He is currently an assistant professor at the Machine Vision Laboratory at
910 the Faculty of Electrical Engineering, University of Ljubljana. His research interests lie in
911 image processing, pattern recognition, embedded design, combinatorial optimization and
912 sequence analysis.

913



914 Janez Perš received B.Sc., M.Sc. and Ph.D. degrees in Electrical engineering at the
915 Faculty of Electrical Engineering, University of Ljubljana, in 1998, 2001 and 2004, re-
916 spectively. He is currently an assistant professor at the Machine Vision Laboratory at
917 the Faculty of Electrical Engineering, University of Ljubljana. His research interests lie
918 in image-sequence processing, object tracking, human-motion analysis, dynamic-motion-
919 based biometry, and in autonomous and distributed systems.

920




921 Rok Mandeljc received his B.Sc. degree in electrical engineering from the Faculty of
922 Electrical Engineering, University of Ljubljana, Slovenia, in 2010. He is currently pursuing
923 the Ph.D. degree as Junior Researcher at Machine Vision Laboratory at the Faculty of
924 Electrical Engineering, University of Ljubljana. His research interests include computer
925 vision, image processing and data fusion in context of multi-view and multi-modal person
926 localization and tracking.

927



928 Vildana Sulić Kenk received her B.Sc. and Ph.D. degrees from the Faculty of Elec-
929 trical Engineering of the University of Ljubljana, Slovenia in 2006 and 2011, respectively.
930 Currently, she is a researcher at the Machine Vision Laboratory at the Faculty of Electri-
931 cal Engineering, with her interests in computer vision, image processing, human-motion
932 analysis and visual-sensor networks.

933



934 Stanislav Kovačič is a professor, a vice dean and the head of the Laboratory for
935 Machine Vision at the Faculty of Electrical Engineering, University of Ljubljana. He was
936 a visiting researcher in the GRASP Laboratory at the University of Pennsylvania, the
937 Technische Fakultät der Friedrich–Alexander–Universität in Erlangen, and at the Faculty
938 of Electrical Engineering and Computing, University of Zagreb. His research is focused on
939 various aspects of image and video analysis with applications in medicine, industry and
940 sports.