

Power-Efficient Approximate Bilateral Filter for Image Denoising on FPGAs Using Flip-Flop-Based Clock Gating

V. Surya¹, Dr. C. Subhas²

¹PG Scholar, Department of ECE, JNTUA College of Engineering, India

²Adjunct Professor, Department of ECE, JNTUA College of Engineering, India

Gmail : vallurusurya95@gmail.com, csubhas.ece@jntua.ac.in

Abstract

This paper presents the hardware design of a power-efficient bilateral filter for image denoising on Field-Programmable Gate Arrays (FPGAs). An approximate-computing kernel strategy is combined with a novel flip-flop-based clock-gating scheme to substantially reduce dynamic power while preserving denoising quality. Spatial-domain and intensity-range kernel coefficients are represented as small unsigned integers, removing costly exponential evaluations and multiplier-heavy arithmetic. A D-flip-flop-synchronized AND-gated clock disables register toggling in the filter datapath whenever a valid pixel window is unavailable, yielding a glitch-free gated clock. Implemented on a Xilinx Zynq-7000 device for a 5×5 kernel and verified with Xilinx Vivado and MATLAB 2020a, the proposed design reduces total on-chip power by approximately 51.3% ($14.863 \text{ W} \rightarrow 7.238 \text{ W}$) and Slice-LUT usage by 18.2% relative to an equivalent design without clock gating, at the cost of a 3.6% increase in critical-path delay. PSNR and SSIM quality metrics remain within -4% and -14% of an accurate software bilateral filter, confirming that the power reduction is obtained without compromising image quality. The results support the suitability of the proposed methodology for real-time, energy-constrained image-denoising applications envisioned under Industry 5.0.

Keywords

Approximate computing, bilateral filtering, FPGA-based architectures, image denoising, low-power design, flip-flop-based clock gating.

I. INTRODUCTION

The transition toward Industry 5.0 is redefining manufacturing by emphasizing sustainability, resilience, and closer collaboration between human operators and intelligent machines. Enabling technologies such as Artificial Intelligence, Edge Computing, Virtual Reality, the Internet of Everything, and Collaborative Robots depend heavily on fast and dependable computer-vision pipelines. Digital image processing sits at the core of these pipelines, supporting human-machine interaction, continuous process monitoring, and real-time decision making on the factory floor.

A persistent obstacle to reliable vision-based decision making is sensor noise. Image sensors and their associated acquisition electronics inevitably corrupt captured frames, and the resulting noise level is rarely known in advance. Consequently, denoising algorithms that can operate effectively across a range of unknown noise levels, and that can do so within tight energy and silicon-area budgets, are of considerable practical importance.

Total variation regularization and bilateral filtering are two of the most widely adopted denoising strategies. Total variation methods preserve edges well but are prone to staircasing and blocky artifacts. Bilateral filtering instead performs a nonlinear weighted average that jointly considers the spatial proximity and the intensity similarity of neighboring pixels, which tends to suppress noise while retaining fine structural detail. The price paid for this quality is computational: the exponential spatial and range kernels involved in bilateral filtering are expensive to evaluate pixel-by-pixel, and a direct hardware mapping of the filter therefore tends to consume considerable logic resources, dynamic power, and clock cycles.

Several FPGA-oriented approximation strategies have been proposed to make bilateral filtering hardware-friendly, typically by replacing the exponential kernels with integer-valued, look-up-table-based approximations that avoid costly multipliers and enable pixel-parallel evaluation. These approaches substantially reduce logic-resource counts and improve throughput, but the great majority of them are optimized for area and speed and give comparatively little attention to dynamic power dissipation, even though switching activity in wide, always-clocked register arrays is a dominant contributor to on-chip power in image-processing pipelines.

This work addresses that gap. It presents the hardware design of a power-efficient, run-time reconfigurable bilateral filter for image denoising, built around an approximate-computing kernel strategy and, as the principal contribution, a flip-flop-based clock-gating scheme applied to the filter datapath. Clock gating disables register toggling in the computation path whenever a valid pixel window is not present, cutting dynamic power without touching the filtering algorithm, the approximation accuracy, or the achievable frame rate. The proposed methodology is implemented and verified on a Xilinx Zynq-7000 FPGA using Vivado, with denoising quality cross-checked in MATLAB.

The main contributions of this paper are summarized as follows:

- An approximate bilateral-filter datapath that represents both the spatial and intensity-range kernel coefficients as small unsigned integers, removing the need for floating-point exponential evaluation and multiplier-heavy arithmetic.
- A flip-flop-based clock-gating architecture, synchronized to the system clock through a D flip-flop and an AND gate, that produces a glitch-free gated clock and eliminates unnecessary switching in the filter's register banks during idle cycles.
- An FPGA implementation and characterization on the Zynq-7000 platform showing a total on-chip power reduction of approximately 51% and an 18.2% reduction in Slice-LUT usage relative to an equivalent design without clock gating, at the cost of a small 3.6% increase in critical-path delay.
- A quality evaluation confirming that the introduced power-saving mechanism leaves the denoising accuracy of the approximate filter unaffected, with PSNR and SSIM remaining within the bounds already established for the underlying approximation strategy.

The remainder of the paper is organized as follows. Section II reviews related bilateral-filtering approximation strategies. Section III recalls the mathematical formulation of bilateral filtering. Section IV describes the proposed kernel-approximation technique. Section V presents the proposed methodology and hardware architecture, including the flip-flop-based clock-gating scheme. Section VI reports implementation and quality results. Section VII concludes the paper.

II. RELATED WORK

Bilateral filtering was originally introduced by Tomasi and Manduchi as an edge-preserving smoothing operator for grayscale and color images, combining a spatial Gaussian kernel with a range kernel defined over intensity differences. Because the direct evaluation of the exponential kernels is costly, a large body of work has since focused on accelerating the filter for real-time and embedded use. Gunturk proposed a fast bilateral filter supporting arbitrary domain and range kernels; Chaudhury, Sage and Unser derived a constant-time formulation using trigonometric range kernels; Paris and Durand introduced a signal-processing view of the filter that enables a fast approximation through downsampling in a higher-dimensional space; and the bilateral-grid formulation of Chen, Paris and Durand further generalized this idea for real-time, edge-aware graphics processing.

On the hardware side, dedicated FPGA architectures have exploited the separability and symmetry of the spatial kernel to group and pipeline pixel processing, approximated the range kernel through low-order polynomial or Taylor expansions, or replaced the conventional exponential kernels with noise-aware, look-up-table-driven piecewise functions to keep logic-resource usage low while remaining adaptable to different noise levels. These architectures report substantial gains in throughput and area relative to a direct, unapproximated implementation, and some support run-time adaptation of the range kernel to the estimated noise level without requiring device reconfiguration.

Despite these advances, most published FPGA bilateral-filter designs are optimized primarily for resource count and frame rate; dynamic power reduction is seldom treated as a first-class design objective, even though the filter's register-heavy, always-active datapath is a significant power contributor in continuous video-processing pipelines. Clock gating is a long-established low-power technique in digital design, but its systematic application to an approximate, run-time-adaptable bilateral-filter datapath has received comparatively little attention. The present work builds on the integer-approximated kernel strategy used in the above architectures and complements it with a flip-flop-based clock-gating scheme purposely inserted into the computation path, so that the resource and power benefits of approximate computing and the switching-activity benefits of clock gating are obtained together.

III. BILATERAL FILTER FUNDAMENTALS

Bilateral filtering denoises a digital image by replacing each pixel with a weighted average of its neighbors, where the weight of a neighboring pixel depends jointly on how close it is in position (the spatial domain) and how similar it is in intensity (the range domain). For a generic pixel $Im(x,y)$ and a $k \times k$ window Ω centered on it, the filtered output $Om(x,y)$ is given by

$$Om(x,y) = \sum_{(i,j) \in \Omega} [Im(i,j) \cdot Ws(i,j) \cdot Wr(i,j)] / \sum_{(i,j) \in \Omega} [Ws(i,j) \cdot Wr(i,j)] \quad (1)$$

where $Ws(i,j)$ and $Wr(i,j)$ are the spatial-domain and intensity-range weighting functions, respectively:

$$Ws(i,j) = \exp(-ED(i,j)^2 / 2\sigma_s^2) \quad (2a)$$

$$Wr(i,j) = \exp(-\Delta I(i,j)^2 / 2\sigma_r^2) \quad (2b)$$

Here $ED(i,j)$ is the Euclidean distance between the pixel positions (x,y) and (i,j) , while $\Delta I(i,j) = |Im(x,y) - Im(i,j)|$ is the corresponding intensity difference. The parameters σ_s and σ_r control how quickly the spatial and range weights decay and are conventionally chosen so that $ED(i,j) \leq 3\sigma_s$ and $\sigma_r = 3\sigma_{noise}$, with σ_{noise} being the estimated noise standard deviation. Because both Ws and Wr are exponential functions evaluated for every pixel in every window, a direct hardware implementation of (1)–(2) is arithmetically expensive; this motivates the approximate kernel strategy described next.

IV. PROPOSED APPROXIMATION TECHNIQUE

A. Approximation of the Spatial-Domain Kernel

To avoid the run-time evaluation of the exponential in (2a), each spatial coefficient $Ws(i,j)$ is pre-computed once, scaled by a factor of 64, and rounded to the nearest 7-bit unsigned integer. The scale factor of 64 is chosen so that the largest coefficient maps to a convenient power-of-two value while keeping the quantization error small. The resulting coefficients are further constrained to have no more than three asserted bits in their binary representation, so that every multiplication by a spatial coefficient reduces to a small number of shift-and-add operations rather than a full multiplier. Because the spatial kernel depends only on the fixed window geometry, this pre-computation is performed once at design time for each supported kernel size.

B. Approximation of the Intensity-Range Kernel

The intensity-range coefficient $Wr(i,j)$ is more difficult to pre-compute because it depends on the pixel data itself through $\Delta I(i,j)$, and on σ_r , which depends on the noise level. The proposed design addresses this with a lightweight piecewise approximation: the intensity difference $\Delta I(i,j)$ is mapped to one of fourteen discrete output levels AW_r , each again representable with no more than four asserted bits, using a small set of comparator-driven thresholds rather than an exponential evaluation. Because the mapping is expressed purely in terms of comparisons against $\Delta I(i,j)$, it can be implemented as a compact look-up table addressed directly by the intensity difference.

To let the same pre-computed threshold table serve a range of noise levels without hardware reconfiguration, a run-time scale factor $sc\sigma$ is introduced such that $\sigma_r = \sigma_{rfix} / sc\sigma$, where σ_{rfix} is the fixed noise level for which the thresholds were originally derived. Scaling the intensity difference by $sc\sigma$ before it addresses the look-up table is mathematically equivalent to evaluating the range kernel at the desired σ_r . With $\sigma_{rfix} = 60$ and $sc\sigma$ represented as a 4-bit fixed-point value ranging from 0.25 to 3.75, the architecture can track effective σ_r values from roughly 16 up to 240, giving the filter run-time adaptability to changing noise conditions with a single, unchanging hardware datapath.

V. PROPOSED METHODOLOGY AND HARDWARE ARCHITECTURE

The proposed design methodology couples the approximate kernel strategy of Section IV with a dedicated low-power datapath, and is organized into four stages: (i) construction of a resource-efficient base filtering architecture; (ii) insertion of a flip-flop-based clock-gating layer around the register-heavy portions of that architecture; (iii) RTL implementation and functional verification; and (iv) FPGA synthesis, implementation, and power characterization. Each stage is detailed below.

A. System-Level Design Flow

At the system level, an input image is first converted to a raw pixel-intensity stream, which is fed to the low-power bilateral-filter core together with the flip-flop-based clock-gating logic that governs its register activity. The core produces a filtered pixel stream that is reassembled into the denoised output image. Fig. 1 summarizes this data flow, highlighting that clock gating acts as a control layer wrapped around the filtering datapath rather than as a modification of the filtering computation itself, so that the numerical behavior of the filter is unaffected by the power-optimization step.




Fig. 1. System-level data flow of the proposed low-power bilateral-filter design.

B. Base Filtering Architecture

The base architecture, shown in Fig. 2, is organized into three cooperating modules. The BUFFER module accepts the incoming $n \times m$ image in raster-scan order and reconstructs the $k \times k$ processing window Ω around the current center pixel using $k-1$ line FIFOs and a bank of $k \times k$ registers, so that a complete neighborhood is available every clock cycle without re-reading pixels from external memory.

The COMPUTATION module contains $k \times k-1$ Processing Elements (PEs) operating in parallel, one for each neighbor of the center pixel. Each PE computes the absolute intensity difference between a neighboring pixel and the center pixel and uses this difference to address a small LUT-based memory that returns the corresponding pre-computed approximate range coefficient $AWr(i,j)$, as derived in Section IV-B.

Finally, the MACs & Divider block combines the $AWr(i,j)$ coefficients produced by the PEs with the pre-computed approximate spatial coefficients $AWs(i,j)$. Because the spatial coefficients are constant powers-of-two combinations at design time, the required multiplications reduce to shifts, and the resulting per-neighbor products are summed through adder trees to form the running numerator and denominator of (1). A non-restoring division stage then produces the final denoised pixel $Om(x,y)$.




FIGURE: Top-level architecture of the BASE design.

Fig. 2. Top-level architecture of the base approximate bilateral-filter design.

C. Proposed Flip-Flop-Based Clock-Gating Scheme

The base architecture of Fig. 2 keeps its FIFOs, window registers, PE registers and accumulator registers clocked on every cycle, regardless of whether a valid pixel window is currently available — for example, during the initial FIFO fill period at the start of every row, or in any cycle where the pipeline is stalled. This continuous toggling of a wide register bank is a major source of dynamic power in the design. To remove it without altering the filtering computation, this work introduces a flip-flop-based clock-gating cell, shown in Fig. 3, inserted immediately ahead of the clock input of the datapath registers.

The cell operates as follows. An enable signal EN, asserted only when a genuine pixel window is present at the input of the COMPUTATION module, is first re-timed by a D flip-flop (FF) clocked on the same edge as the datapath, producing a synchronized enable that changes state only immediately after an active clock edge. This synchronized enable is then combined with the free-running clock CLK through a two-input AND gate

to produce the gated clock CLK_GATED. Because the enable signal is captured by a flip-flop before being AND-ed with the clock — rather than being gated with a transparent latch, as in conventional latch-based clock gating — the resulting CLK_GATED signal is inherently free of the mid-cycle glitches that latch-based gating can introduce, since the enable can only change synchronously with the clock edge.




Fig. 3. Flip flop based clock gating technique

Fig. 3. Proposed flip-flop-based clock-gating cell used to drive the datapath registers of the bilateral-filter core.

When EN is de-asserted — for instance while a FIFO is still filling, or between valid frames — CLK_GATED is held permanently low, so every register downstream of the gating cell stops toggling entirely and its dynamic power contribution drops to essentially zero for the duration of the idle interval. When EN is asserted, CLK_GATED reproduces the original clock with no added skew beyond the small, well-controlled delay of the AND gate, so the filtering pipeline continues to operate exactly as in the base design during active cycles. Because the MAC-and-divider datapath of the bilateral filter is only required once a full $k \times k$ window has been assembled, gating its registers in this way targets precisely the cycles in which the register bank would otherwise be toggling without producing useful output.

This clock-gating layer is purely a control-path addition: it does not modify the approximate spatial or range kernels, the adder-tree structure, or the non-restoring divider described in Section V-B, and therefore leaves the numerical output of the filter, and hence its PSNR/SSIM behavior, unchanged relative to the base design.

D. Implementation and Verification Methodology

The proposed methodology was carried out in the following sequence. First, the base approximate bilateral-filter datapath (Section V-B) was described in RTL and functionally verified against a MATLAB 2020a software reference to confirm that the integer-approximated kernels reproduce the expected denoising behavior. Second, the flip-flop-based clock-gating cell (Section V-C) was integrated into the RTL, driving the clock input of the BUFFER, PE and MACs & Divider register banks, and the combined design was re-verified in simulation to confirm bit-exact agreement with the base design's output while exercising the enable/idle transitions of the gating cell. Third, both the base design and the clock-gated design were synthesized and implemented for the Xilinx Zynq-7000 device family using Xilinx Vivado, targeting a 5×5 kernel with $\sigma_s = 1$, so that resource utilization, timing and power could be compared on an

VI. RESULTS AND DISCUSSION

Post-implementation simulation of the proposed design in Vivado confirmed correct end-to-end operation: the filter correctly asserted its done flag and produced the expected filtered pixel stream after processing a complete 256×256 test image, with the clock-gating enable/idle transitions observed to behave as intended — CLK_GATED toggled only while a valid pixel window was present at the COMPUTATION module and remained low during FIFO-fill and idle intervals.


[illegible]

Fig. 5. RTL schematic of the proposed clock-gated bilateral-filter design (Vivado).

Table I compares the base design (no clock gating) with the proposed clock-gated design, both implemented for a 5×5 kernel on the Zynq-7000 device. The dominant result is the reduction in total on-chip power, from 14.863 W in the base design to 7.238 W in the proposed design, a decrease of approximately 51.3%, obtained purely through the clock-gating layer of Section V-C without any change to the filtering algorithm or to image quality.

Parameter	Base Design	Proposed Design	Change
Slice LUTs	681	557	-18.2%
Slice Registers	8	9	+1 (FF)
Bonded IOBs	114	115	~same
BUFGCTRL	1	1	same
Total Power (W)	14.863	7.238	-51.3%
Crit. Path Delay (ns)	7.900	8.182	+3.6%
Junction Temp. (°C)	~55.0	38.5	Improved

TABLE I. Comparison of Base and Proposed (Clock-Gated) Designs (Zynq-7000, 5×5 kernel).

Alongside the power reduction, Slice-LUT usage falls from 681 to 557, an 18.2% reduction, since the clock-gated design no longer requires as much always-active enabling logic distributed through the datapath. Slice-Register count increases only marginally, from 8 to 9, reflecting the single additional flip-flop used to synchronize the enable signal in the gating cell; the number of Bonded IOBs and the BUFCTRL count used for global clock distribution remain essentially unchanged. The critical-path delay increases slightly, from 7.9 ns to 8.182 ns (a 3.6% increase), which is attributable to the AND gate inserted in the clock path; this increase is small enough to remain comfortably within typical timing margins and does not compromise real-time operation. As a secondary benefit, the reduced switching activity also lowers the estimated junction temperature from approximately 55°C to 38.5°C, indicating improved thermal margin for sustained operation.




Fig. 6. Vivado resource-utilization, timing and power report for the proposed clock-gated design.

Because the clock-gating layer only affects when the datapath registers are clocked, and not the values they compute, the denoising quality of the proposed design tracks that of the underlying approximate kernel strategy of Section IV. On the USC-SIPI Miscellaneous dataset, across noise standard deviations σ_{noise} ranging from 5 to 60 and kernel sizes from 3×3 to 21×21 , the approximate filter's PSNR and SSIM remain within about -4% and -14% , respectively, of an accurate MATLAB software bilateral filter used as the reference. On a representative test image corrupted with additive Gaussian noise of $\sigma_{\text{noise}} \approx 30$, the proposed filter visibly suppresses small-scale speckle noise in flat background regions while preserving the edges and fine structure of the dominant foreground content.

consistent with the quantitative PSNR/SSIM results, as illustrated in Fig. 7.




Fig. 7. Input image corrupted with Gaussian noise ($\sigma_{\text{noise}} \approx 30$) and corresponding output image after proposed bilateral-filter denoising.

D. Discussion

Taken together, these results indicate that flip-flop-based clock gating is an effective and low-risk complement to approximate-kernel bilateral filtering on FPGAs: it delivers roughly a two-fold reduction in total on-chip power and a meaningful reduction in LUT usage, at the cost of a small, easily accommodated timing overhead and no measurable loss of denoising quality. Because the gating cell is inserted purely in the clock path and re-uses the existing enable signal that already indicates window validity in the base architecture, it can be adopted with limited modification to existing approximate bilateral-filter datapaths, including run-time-reconfigurable variants that adapt or through the scale factor scor described in Section IV-B.

VII. CONCLUSION

This paper presented the design and FPGA implementation of a power-efficient, approximate bilateral filter for image denoising that combines an integer-approximated spatial/range kernel strategy with a novel flip-flop-based clock-gating scheme applied to the filter's register-heavy datapath. The proposed clock-gating cell synchronizes its enable signal with the system clock through a D flip-flop before AND-ing it with the clock, yielding a glitch-free gated clock that suppresses register toggling during idle cycles without altering the filtering computation. Implemented on a Xilinx Zynq-7000 device for a 5×5 kernel, the proposed design reduces total on-chip power by approximately 51.3% (from 14.863 W to 7.238 W) and Slice-LUT usage by 18.2% relative to an equivalent design without clock gating, while incurring only a 3.6% increase in critical-path delay and leaving PSNR/SSIM denoising quality unchanged. These results indicate that flip-flop-based clock gating is a practical, low-overhead technique for reducing the power footprint of run-time-adaptable bilateral-filter hardware, making the proposed design well suited to energy-constrained, real-time image-denoising applications such as industrial automation, medical imaging, autonomous vehicles and smart-city IoT systems envisioned under Industry 5.0.

VIII. AUTHOR'S CONTRIBUTION

Conceptualization: V. Surya. Methodology: V. Surya. Investigation: V. Surya. Writing – Original Draft: V. Surya.

Supervision: Dr. C. Subhas. Approval of the final text: V. Surya and Dr. C. Subhas.

IX. ACKNOWLEDGMENTS

The authors thank Jawaharlal Nehru Technological University Ananthapur College of Engineering (Autonomous) for infrastructure support, and Dr. C. Subhas (Adjunct Professor, Department of ECE, JNTUACEA) for guidance throughout this work.

X. REFERENCES

- [1] P. K. R. Maddikunta, Q.-V. Pham, B. Prabadevi, N. Deepa, and K. Dev, "Industry 5.0: A survey on enabling technologies and potential applications," *J. Ind. Inf. Integr.*, vol. 26, pp. 1–31, Mar. 2022, doi: 10.1016/j.jii.2021.100257.
- [2] F. Spagnolo, P. Corsonello, F. Frustaci, and S. Perri, "Design of approximate bilateral filters for image denoising on FPGAs," *IEEE Access*, vol. 11, pp. 1990–2000, Jan. 2023, doi: 10.1109/ACCESS.2022.3233921.
- [3] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. 6th Int. Conf. Comput. Vis.*, Jan. 1998, pp. 839–845, doi: 10.1109/ICCV.1998.710815.
- [4] S. Prasath and D. N. Thanh, "Structure tensor adaptive total variation for image restoration," *Turkish J. Electr. Eng. Comput. Sci.*, vol. 27, pp. 1147–1156, Mar. 2019, doi: 10.3906/elk-1802-76.
- [5] B. K. Gunturk, "Fast bilateral filter with arbitrary range and domain kernels," *IEEE Trans. Image Process.*, vol. 20, no. 9, pp. 2690–2696, Sep. 2011.
- [6] B.-H. Chen, Y.-S. Tseng, and J.-L. Yin, "Gaussian-adaptive bilateral filter," *IEEE Signal Process. Lett.*, vol. 27, pp. 1670–1674, 2020.
- [7] K. N. Chaudhury, D. Sage, and M. Unser, "Fast $O(1)$ bilateral filtering using trigonometric range kernels," *IEEE Trans. Image Process.*, vol. 20, no. 12, pp. 3376–3382, 2011.
- [8] P. Choudhury and J. Tumblin, "The trilateral filter for high contrast images and meshes," in *Eurographics Symp. Rendering*, 2003.
- [9] S. Paris and F. Durand, "A fast approximation of the bilateral filter using a signal processing approach," *Int. J. Comput. Vis.*, vol. 81, no. 1, pp. 24–52, 2009.
- [10] T. Q. Pham and L. J. van Vliet, "Separable bilateral filtering for fast video preprocessing," in *IEEE Int. Conf. Multimedia and Expo*, 2005.
- [11] A. Adams, J. Baek, and M. A. Davis, "Fast high-dimensional filtering using the permutohedral lattice," *Comput. Graph. Forum*, vol. 29, no. 2, pp. 753–762, 2010.
- [12] J. Chen, S. Paris, and F. Durand, "Real-time edge-aware image processing with the bilateral grid," *ACM Trans. Graph.*, vol. 26, no. 3, 2007.
- [13] M. Elad, "On the origin of the bilateral filter and ways to improve it," *IEEE Trans. Image Process.*, vol. 11, no. 10, pp. 1141–1151, 2002.
- [14] V. Aurich and J. Weule, "Non-linear Gaussian filters performing edge preserving diffusion," in *DAGM Symp. Pattern Recognit.*, 1995.
- [15] A. Hore and D. Ziou, "Image quality metrics: PSNR vs. SSIM," in *Proc. Int. Conf. Pattern Recognit.*, 2010, pp. 2366–2369, doi: 10.1109/ICPR.2010.579.