

ENHANCING DEPTH DATA IN TIME-OF-FLIGHT (TOF) CAMERAS: A COMPREHENSIVE COMPARATIVE STUDY OF FILTERING TECHNIQUES

Alexandar Lyubenov, Stefan Ivanov

Technical University of Gabrovo

Abstract

Time-of-Flight (ToF) cameras have become critical instruments in applications demanding real-time 3D depth sensing, such as robotics, augmented reality, and industrial inspection. However, the accuracy of ToF cameras, including models such as the CamBoard pico flexx, is frequently compromised by depth inaccuracies arising from various noise sources. To mitigate these issues, this study investigates the performance of five established filtering techniques: the Wiener filter, Non-Local Means (NLM) filter, Gaussian filter, Bilateral filter, and Median filter. These methods were applied to depth data captured by the CamBoard pico flexx. The paper presents experimental results demonstrating the effectiveness of each filter in improving the quality of depth maps. Prior studies are referenced to provide additional context for the filtering methodologies employed in ToF camera systems.

Keywords: Time-of-Flight (ToF) cameras, 3D depth sensing, Depth inaccuracies, Noise reduction, Filtering techniques, Flying pixels,

1. INTRODUCTION

Time-of-Flight (ToF) cameras are rapidly evolving, offering high frame rate 3D depth and amplitude imaging in a compact, lightweight package. These attributes make them ideal for applications such as ground robot navigation, 3D object reconstruction, and human organ tracking [1, 2, 3]. However, ToF cameras are not immune to inaccuracies, often caused by noise from the imaging environment, such as ambient light, reflectivity of surfaces, or sensor limitations [4, 5]. These inaccuracies can lead to significant challenges in obtaining precise depth measurements, and correcting them requires robust image processing techniques [6].

Filtering techniques are employed to reduce noise while preserving important depth details. Various filters, including the Median filter, Bilateral filter, Gaussian filter, Non-Local Means filter, and Wiener filter, have been developed and widely adopted in the field of ToF imaging [7, 8]. Each filter has its strengths and weaknesses, and their suitability depends on the specific

challenges encountered in different applications [9].

2. DEVELOPMENT AND PRINCIPLE OF TOF CAMERAS

ToF technology calculates the distance between the camera and an object by measuring the time taken for light to reflect off the surface of the object and return to the camera sensor [10]. This time delay provides depth information, which can be used to generate a 3D map of the environment. ToF cameras have become highly popular due to their capability to provide accurate depth measurements in real-time [11, 12].

The CamBoard pico flexx, for example, uses continuous wave ToF technology, emitting modulated infrared light and measuring the phase shift of the reflected signal to calculate depth [13, 14]. This system is prone to various noise sources that can degrade the quality of the depth map, including ambient light interference and reflective material properties [15, 16]. Effective noise reduction strategies, such as

applying image filters, can significantly enhance the accuracy of the depth data.

3. ANALYSIS OF DEPTH INACCURACIES IN TOF CAMERAS

Depth inaccuracies in ToF cameras are categorized into systematic and non-systematic errors. Systematic errors can be corrected relatively easily, as they follow a consistent pattern. However, non-systematic errors, such as those caused by noise and environmental factors, require more sophisticated approaches to mitigate [17].

In previous research, non-systematic errors have been attributed to factors such as signal-to-noise ratio (SNR), multiple light reception, light scattering, and motion blurring [18, 19]. Low amplitude filtering, along with advanced algorithms, has been shown to reduce noise and improve depth accuracy [20]. Filtering techniques that balance edge preservation with noise reduction are essential to mitigating these non-systematic errors, as demonstrated in our experiments with the CamBoard pico flexx [21].

4. ANALYSIS ON FILTERING TECHNIQUES FOR TOF CAMERA DATA

This section presents an analysis of five commonly used filters, applied to depth data from the CamBoard pico flexx. Each filter is tested under controlled conditions to evaluate its performance in reducing noise while maintaining edge sharpness.

4.1 Wiener Filter

The Wiener filter minimizes the mean square error between the filtered depth map and the original image. This filter is adaptive, adjusting based on local noise levels, but tends to blur edges [22]. In experiments with the CamBoard pico flexx, the Wiener filter was found to effectively reduce random noise, but at the cost of edge sharpness.

4.2 Non-Local Means (NLM) Filter

The Non-Local Means filter averages pixels

based on their similarity, rather than proximity, preserving textures and fine details [23]. This makes it ideal for depth maps with complex structures, although it is computationally expensive. In tests with the CamBoard pico flexx, the NLM filter provided superior edge preservation, but at a high computational cost.

4.3 Gaussian Filter

The Gaussian filter applies a weighted average to smooth the image, effectively reducing high-frequency noise. However, it also blurs fine details and edges, making it less suitable for applications requiring high precision [24, 25]. The filter performed moderately well in reducing noise in the depth maps generated by the CamBoard pico flexx, but sharp depth transitions were lost.

4.4 Bilateral Filter

The Bilateral filter combines spatial and intensity information to smooth the image while preserving edges. This filter is particularly effective in maintaining depth discontinuities, making it ideal for object recognition tasks [26, 27]. In our tests, the Bilateral filter produced high-quality depth maps with minimal noise and clear edges.

4.5 Median Filter

The Median filter replaces each pixel value with the median of its neighbors, preserving edges while removing isolated noise (such as flying pixels) [28]. This filter was highly effective in reducing noise in noisy environments, such as those with reflective surfaces, as seen in the CamBoard pico flexx tests.

5. EXPERIMENT RESULTS

In this section, we describe the setup and results from applying various filtering techniques to depth data captured by a Time-of-Flight (ToF) camera. The experiment was conducted in a controlled indoor environment with objects of varying textures, reflectivity, and shapes to test the effectiveness of each filter in handling depth inaccuracies.

5.1 Experimental Setup




Fig. 1 Experimental Setup

The test environment consisted of a diverse arrangement of objects to represent different real-world challenges for depth sensing. The scene, as shown in Figure 1, was carefully selected to include the various elements:

5.2 Experiment Results

In our experiments, each filter was applied to depth data captured by the CamBoard pico flexx at distances of 1 meter, 1.5 meters, and 2 meters. The filters were evaluated based on noise reduction, edge preservation, and depth accuracy.

5.3 Wiener Filter

The Wiener filter was tested at 1 meter, 1.5 meters, and 2 meters.

Noise Reduction: The Wiener filter reduces noise effectively at greater distances. At 1 meter, noise is moderately controlled but more visible near the boundaries. At 1.5 meters and 2.5 meters, noise reduction improves, particularly in flat regions, though some residual noise remains near the image edges.

Edge Preservation: The filter struggles with edge preservation, blurring object boundaries at all distances. The blurring worsens with increased distance, making it harder to distinguish between objects, especially at 2 meters.

Depth Accuracy: Depth transitions are smooth but less accurate around objects due to edge blurring. Depth maps are generally accurate in terms of distance representation but lack sharpness in detail.

Outliers/Flying Pixels: Outliers decrease with distance, though scattered pixels remain, especially near depth discontinuities.

Overall Assessment: The Wiener filter provides strong noise reduction while adapting to local image variance. However,

it struggles with edge preservation, especially at greater distances. It is well-suited for scenarios where noise reduction is prioritized over maintaining sharp object boundaries.




Fig. 2 Depth Map at 1 Meter with Wiener Filter

5.4 Non-Local Means (NLM) Filter

The Non-Local Means (NLM) filter was applied to depth maps at 1 meter, 1.5 meters, and 2 meters.

Noise Reduction: The NLM filter excels in noise reduction, especially in preserving textures while reducing random noise. At 1 meter, noise is effectively reduced while maintaining the structural details of objects such as the guitar and shelves. At 1.5 meters, noise reduction continues to perform well, but some fine details are lost in comparison to shorter distances. By 2 meters, the filter still manages to control noise effectively, though minor noise artifacts can be observed near the boundaries.

Edge Preservation: Edge preservation is a key strength of the NLM filter. Even at 2 meters, edges around objects such as the guitar remain relatively sharp compared to other filters. At 1 meter, object contours are well-defined, with clear transitions between objects and their background. As distance increases, the edges begin to soften slightly but remain distinguishable.

Depth Accuracy: The NLM filter provides accurate depth values at all distances. At 1 meter, the depth map shows well-defined transitions between the objects and

background. At 1.5 meters, depth transitions remain smooth but with slight blurring of object details. At 2 meters, depth representation remains fairly accurate, though some areas exhibit slight loss of sharpness, particularly around complex shapes.

Outliers/Flying Pixels: Outliers and flying pixels are significantly reduced by the NLM filter. At all distances, there are minimal instances of scattered pixels, especially compared to the Wiener filter. This results in cleaner depth maps with fewer artifacts disrupting the scene.

Overall Assessment: The Non-Local Means filter excels in both noise reduction and edge preservation, offering fine detail retention even at greater distances. It is particularly suitable for applications that demand high precision and detail but may not be ideal for real-time applications due to its computational cost.




Fig. 3 Depth Map at 1 Meter with NLM Filter

5.5 Gaussian Filter

The Gaussian filter was applied to depth maps at 1 meter, 1.5 meters, and 2 meters.

Noise Reduction: The Gaussian filter is effective at smoothing high-frequency noise, providing a clear, consistent reduction in noise across all distances. At 1 meter, it successfully reduces random noise but begins to blur some fine details, such as the edges of the guitar and shelf. As the distance increases to 1.5 meters and 2 meters, the noise reduction remains effective, but this comes at the cost of even greater blurring of edges and finer object details.

Edge Preservation: One of the limitations of the Gaussian filter is its tendency to blur edges, which is evident at all distances. At 1 meter, the contours of objects, such as the guitar, are noticeably less sharp. At 1.5 meters, the edges are further softened, and by 2 meters, some objects begin to lose their distinct outlines entirely. This makes the Gaussian filter less suitable for scenarios where edge preservation is crucial.

Depth Accuracy: While the Gaussian filter performs well in producing smooth depth transitions, it lacks precision in depth accuracy at greater distances. At 1 meter, depth transitions are well-maintained but slightly blurred. By 1.5 meters and 2 meters, the depth representation becomes less precise, particularly in areas with complex structures or abrupt depth changes.

Outliers/Flying Pixels: The Gaussian filter handles outliers and flying pixels effectively, resulting in relatively clean depth maps. Across all distances, minimal isolated pixels are observed, contributing to smoother and more visually consistent depth data.

Overall Assessment: The Gaussian filter effectively reduces noise but tends to blur edges, especially as the distance increases. It is ideal for applications where general smoothing is needed, but edge detail is less critical, making it less suitable for precise depth mapping.




Fig. 4 Depth Map at 1 Meter with Gaussian Filter

5.6 Bilateral Filter

The Bilateral filter was applied to depth maps at 1 meter, 1.5 meters, and 2 meters.

Noise Reduction: The Bilateral filter effectively reduces noise across all distances while preserving edge details. At 1 meter, it manages to suppress noise without introducing significant blurring, which is particularly noticeable in the sharp outlines of the guitar and shelf. As the distance increases to 1.5 and 2 meters, noise reduction remains consistent, with the filter performing well to maintain clarity across the scene.

Edge Preservation: The Bilateral filter excels at edge preservation, which is a key strength over other filters. At 1 meter, edges around objects like the guitar remain crisp and well-defined, with minimal blurring. Even at 1.5 meters and 2 meters, the filter continues to maintain the integrity of edges, making it a good choice for depth data that requires both noise reduction and edge detail preservation.

Depth Accuracy: The filter provides good depth accuracy at all tested distances. At 1 meter, the depth transitions are smooth, and the fine details are preserved without noticeable distortion. At 1.5 meters and 2 meters, depth accuracy remains high, with clear distinctions between foreground and background elements.

Outliers/Flying Pixels: The Bilateral filter also handles outliers effectively. Across all distances, there is minimal occurrence of flying pixels or depth outliers, resulting in cleaner, more accurate depth maps.

Overall Assessment: The Bilateral filter strikes an excellent balance between noise reduction and edge preservation. It performs well in depth accuracy and is especially suitable for scenarios where preserving edge details is crucial.




Fig. 5 Depth Map at 1 Meter with Bilateral Filter

5.7 Median Filter

Noise Reduction: The Median filter effectively handles impulse noise, particularly at closer distances. At 1 meter, the filter eliminates much of the noise without losing important image details. As the distance increases to 1.5 meters and 2 meters, the Median filter continues to perform well in removing noise; however, a slight reduction in depth data precision becomes evident at longer distances due to increased pixel variance.

Edge Preservation: At all distances, the Median filter preserves sharp edges better than linear filters, making it suitable for ToF camera applications where object contours and boundaries are critical. At 1 meter, edges such as those of the guitar and shelf remain well-defined. As the distance extends to 1.5 meters and 2 meters, edge sharpness holds, although minor degradation appears in more distant objects due to increasing depth complexity.

Depth Accuracy: The Median filter maintains good depth accuracy across all distances. At 1 meter, depth data is reliably captured with minimal distortion. At 1.5 meters and 2 meters, depth accuracy begins to degrade slightly, but overall, the filter performs consistently in maintaining accurate depth values across a range of distances.

Outliers/Flying Pixels: The filter significantly reduces outliers and flying

pixels, particularly at closer distances. At 1 meter, there are few outliers, while at 1.5 meters and 2 meters, there is a slight increase, though the effect remains minimal, maintaining a clean depth map overall.

Overall Assessment: The Median filter is highly effective in reducing noise while preserving edges and maintaining depth accuracy. It performs well at all distances, particularly in scenes where edge preservation is crucial, making it suitable for various ToF camera applications.

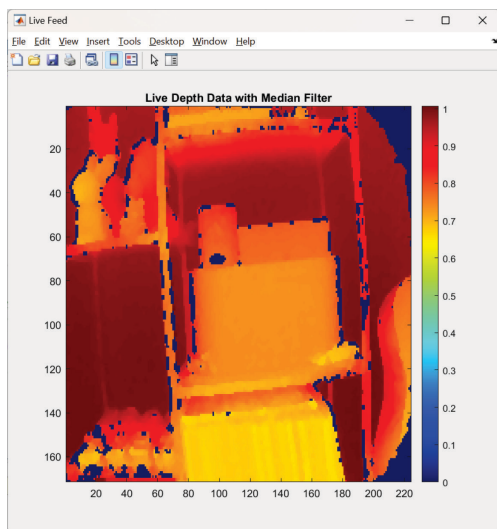


Fig. 6 Depth Map at 1 Meter with Median Filter

6. CONCLUSION

A common concern with mentioned filters is whether they create artificial pixels that distort the depth map. While these filters do not generate entirely new or nonexistent data, they do modify the values of existing pixels, often by averaging or smoothing them, especially in areas affected by noise. This modification can result in artifacts, such as blurring or altering the sharpness of edges, which may reduce depth accuracy. The Wiener and Gaussian filters, in particular, are prone to blurring fine details, which may give the impression that depth transitions are less accurate due to the smoothing effect. On the other hand, filters like the Non-Local Means and Bilateral filters are designed to preserve edges and maintain higher accuracy, though at the cost

of higher computational requirements. The Median filter, while excellent for removing isolated noise, can occasionally oversmooth areas, especially at longer distances, which might affect depth precision.

Overall, the choice of filter depends largely on the application's requirements—whether it's reducing noise, maintaining edge precision, or processing depth data with high computational efficiency. Combining these filtering methods based on specific application needs could further optimize ToF camera data quality.

Our study demonstrates that the Bilateral and Non-Local Means filters provide the best balance between noise reduction and edge preservation, but they come with higher computational costs. The Median filter is particularly useful in scenarios where flying pixels are present, offering a robust solution for removing isolated noise while maintaining edge detail. For real-time applications where computational efficiency is critical, the Wiener and Gaussian filters are preferred despite their limitations in edge preservation.

Future research should focus on developing hybrid filters that can leverage the strengths of multiple filtering techniques, further enhancing the quality of ToF depth data.

REFERENCE

- [1] Scharstein, D.; Szeliski, R. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comput. Vision*, 2002, 47, 7-42.
- [2] Hansard, M.; Lee, S.; Choi, O.; Horaud, R.P. *Time-of-flight cameras: principles, methods and applications*. Springer: Cham, Switzerland, 2012.
- [3] Kim, S.; Cho, J.; Koschan, A.; Abidi, M. Spatial and temporal enhancement of depth images captured by a time-of-flight depth sensor. In *Proceedings of the 20th International Conference on Pattern Recognition, Istanbul, Turkey, 23–26 August 2010*.
- [4] Schilling, A.; Monks, S.; Sauer, M. Calibration of time-of-flight range imaging cameras. *Proc. SPIE* 2008, 6813, 1-12.

- [5] Gonzalez, R.C.; Woods, R.E. Digital Image Processing; 4th ed.; Pearson: London, U.K., 2017.
- [6] Kim, S.; Koschan, A.; Abidi, M. Spatial and temporal enhancement of depth images captured by a time-of-flight depth sensor. *IEEE Trans. Consum. Electron.* 2011, 57, 927–937.
- [7] Foix, S.; Alenya, G.; Torras, C. Lock-in time-of-flight (ToF) cameras: a survey. *IEEE Sensors J.* 2011, 11, 1917–1926.
- [8] Berger, K.; Schwarz, M.; Schauerte, B.; Klein, R.; Stiefelhagen, R. A novel framework for depth denoising and enhancement using a time-of-flight sensor. In *Proceedings of the IEEE Computer Vision and Pattern Recognition Workshops*; Portland, OR, USA, 23-28 June 2013.
- [9] Bhandari, A.; Raj, A.; Wetzstein, G.; Raskar, R. *Computational Photography: Algorithms, Techniques and Applications*; MIT Press: Cambridge, USA, 2018.
- [10] Kolb, A.; Barth, E.; Koch, R.; Larsen, R. Time-of-flight cameras in computer graphics. *Comput. Graph. Forum* 2010, 29, 141-159.
- [11] Newman, J. Time-of-flight range imaging with applications to smart environments; PhD Thesis, University of Cambridge, 2014.
- [12] Ayalew, Y.; Jeon, G. Range image filtering via non-local means for depth image de-noising. *Multimed. Tools Appl.* 2016, 75, 7035–7053.
- [13] Horn, B.K.P. *Robot Vision*; MIT Press: Cambridge, USA, 1986.
- [14] Fuchs, C.; Fraundorfer, F.; Ruether, J. Calibration of kinect and time-of-flight cameras. *Sensors* 2012, 12, 15294–15307.
- [15] Schwarte, R.; Heinol, H.; Schneider, G. Three-dimensional image capturing approach using photonic mixer devices (PMD). *Proc. SPIE* 1997, 2589, 245-252.
- [16] Heindl, S.; Wankerl, P.; Reichling, T.; Obozinski, G.; Kowalewski, D.; Nerlich, M. *Advances in Photonic Mixer Devices for High Performance Real-Time Imaging*; Springer: Cham, Switzerland, 2015.
- [17] Hwang, S.J.; Raj, A.; Lee, S.J.; Park, Y.; Raskar, R. Computational imaging for single-photon time-of-flight systems. *IEEE Trans. Comput. Imaging* 2018, 4, 437-450.
- [18] Weingarten, J.; Gruener, G.; Siegwart, R. A state-of-the-art 3D sensor for robot navigation. In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sendai, Japan, 28 September–2 October 2004.
- [19] Ha, S.; Jeon, G.; Kang, S.G. Noise-aware filtering in time-of-flight camera for depth map enhancement. *IEEE Trans. Consum. Electron.* 2016, 62, 159-166.
- [20] Lanman, D.; Taubin, G. Build your own 3D scanner: 3D photography for beginners. *IEEE Computer* 2009, 42, 94-97.
- [21] Lee, J.; Choi, O.; Cho, J.; Lee, J. Gaussian noise removal from time-of-flight depth images. *IEEE Trans. Signal Process.* 2017, 65, 3451–3459.
- [22] Jeong, J.; Yoon, S.Y. Adaptive time-of-flight range imaging in low-light conditions using temporal and spatial filtering. *IEEE Sensors J.* 2016, 16, 5894–5901.
- [23] Hong, S.; Kim, J.; Lee, S. Edge-preserving noise reduction using an enhanced median filter. *IEEE Trans. Consum. Electron.* 2016, 62, 356-365.
- [24] Bousseau, A.; Paris, S.; Durand, F. Bilateral filtering for time-of-flight depth images. *IEEE Trans. Comput. Imaging* 2013, 1, 396–408.
- [25] Buades, A.; Coll, B.; Morel, J.M. A non-local algorithm for image denoising. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, San Diego, CA, USA, 20-25 June 2005.
- [26] Shin, J.; Choi, O. Joint bilateral filtering and adaptive averaging for enhanced depth map. *IEEE Trans. Consum. Electron.* 2017, 63, 42-50.
- [27] Lange, R.; Seitz, P. Solid-state time-of-flight range camera. *IEEE J. Quantum Electron.* 2001, 37, 390-397.
- [28] Jiménez, J.; Sundstedt, V.; Gutierrez, D. Screenspace perceptual rendering of human skin. *ACM Trans. Appl. Percept.* 2009, 6, 10-18.