OT-2026-000000047

Bin Picking Depth Estimation Using Dual Projector Angle Rejection

Published	September 15, 2026 at 11:48:03 UTC
Permanent URL	https://opentechnical.org/p/OT-2026-000000047
Classification	Sensors, Controls & Machine Vision > Machine Vision > 3D Vision
Document SHA-256	7baba6e2780f29b8bfcafa828b3abd1943700a35d1ef376ded11c10877275d07
Extent	3,021 words, 21 numbered paragraphs, 2 figures

Abstract

This disclosure describes an optical depth measurement system for robotic bin picking applications, designed to mitigate errors caused by specular reflections from object surfaces. Conventional structured light systems can misinterpret specular returns as valid surface points, leading to inaccurate depth maps and failed pick attempts. The described arrangement employs two structured light projectors positioned at distinct angles relative to an imaging sensor. Depth maps generated from each projector are cross-referenced to identify and reject inconsistent data points, which often originate from specular highlights. A minimum density of valid depth points is required before a pick operation, with an automatic exposure adjustment as a fallback to ensure sufficient data acquisition.

1. Technical Field

[0001] The present disclosure relates to machine vision systems, specifically to three-dimensional (3D) depth estimation for robotic manipulation, such as bin picking, employing structured light techniques. More particularly, it addresses methods for improving the robustness of depth sensing in environments containing objects with highly reflective or specular surfaces.

2. Technical Context

[0002] In automated manufacturing and logistics, robotic systems are increasingly deployed for tasks like bin picking, where a robot must identify and retrieve individual objects from a disordered collection. Accurate 3D perception is critical for these applications, enabling the robot to localize objects, determine grasp points, and plan collision-free trajectories. Structured light systems are commonly used for 3D perception, projecting a known pattern onto a scene and observing its deformation with an imaging sensor (14) to calculate depth.

[0003] A significant challenge in structured light depth estimation arises when objects possess highly reflective or specular surfaces. When a structured light pattern from a projector (10, 12) strikes a specular surface, the light is reflected at an angle equal to the angle of incidence. If this reflected light directly enters the imaging sensor (14), it can create a false bright spot or pattern distortion that the depth processing algorithms interpret as a valid surface point, often corresponding to an incorrect depth. This phenomenon, known as specular reflection, leads to erroneous depth maps. A robot attempting to pick an object based on such an erroneous depth map may reach into empty space, collide with other objects, or fail to grasp the intended item, thereby reducing efficiency and potentially causing damage. The problem is exacerbated in industrial settings where parts are often metallic, polished, or coated, presenting numerous specular features. Therefore, a method to reliably distinguish and reject these spurious depth readings is necessary to improve the accuracy and robustness of bin picking systems.

3. System Overview

[0004] The system provides a method for generating robust 3D depth maps of objects, particularly those with specular surfaces, for robotic manipulation. As illustrated in FIG. 1, the core components include two structured light projectors (10, 12), an imaging sensor (14), a depth processor (16), a validation unit (18), and a control unit (20). The two projectors (10, 12) sequentially or simultaneously project structured light patterns onto the scene. The imaging sensor (14) captures images of the scene with the projected patterns. The depth processor (16) processes these images to generate an initial depth map for each projector. The validation unit (18) then cross-references these two depth maps to identify and reject inconsistent depth points, which are highly likely to be specular reflections. The control unit (20) manages the projection sequence, image acquisition, exposure settings, and evaluates the density of valid depth points. If the density of valid points is insufficient for a reliable pick, the control unit (20) can initiate a fallback procedure, such as adjusting exposure settings, to acquire more data. Once a sufficient number of validated depth points are available, the control unit (20) can issue a pick command to a robot system (22).

4. Components

[0005] The system comprises several interconnected components, as depicted in FIG. 1 and FIG. 2.

[0006]

- **Projector 1 (10):** A structured light projector configured to emit a specific light pattern onto the target scene. This pattern can be a series of parallel lines, a grid, a random dot pattern, or other codings. The projector (10) is positioned at a defined angle relative to the imaging sensor (14).
- **Projector 2 (12):** A second structured light projector, identical in type and capability to Projector 1 (10). It is positioned at a different, distinct angle relative to the imaging sensor (14) compared to Projector 1 (10). The angular separation between the two projectors (10, 12) is a critical design parameter.
- **Imaging Sensor (14):** A high-resolution digital camera or similar optical sensor capable of capturing images of the scene illuminated by the structured light patterns. It is typically equipped with a lens system optimized for the working distance and field of view. The imaging sensor (14) is positioned to observe the scene from a specific viewpoint.
- **Depth Processor (16):** A computational module, which may be implemented in hardware (e.g., FPGA, ASIC) or software running on a general-purpose processor. Its function is to analyze the images captured by the imaging sensor (14) and, using triangulation principles based on the known projector (10, 12) and sensor (14) geometry, compute a depth map for each projected pattern. This module generates two distinct depth maps, one corresponding to Projector 1 (10) and one to Projector 2 (12).
- **Validation Unit (18):** A computational module responsible for comparing the two depth maps generated by the depth processor (16). It identifies corresponding points in the two maps and performs a disagreement test based on their estimated depth values. Points where the depth values are inconsistent beyond a predefined tolerance are marked as invalid.

- **Control Unit (20):** The central processing unit of the system. It orchestrates the operation of the projectors (10, 12), the imaging sensor (14), and the depth processor (16). It receives validated depth data from the validation unit (18), evaluates the quality and density of the resultant point cloud, and makes decisions regarding subsequent actions, such as initiating a robot pick or adjusting system parameters. It also manages exposure settings for the imaging sensor (14).
- **Robot System (22):** An external robotic manipulator equipped with an end effector (e.g., gripper, suction cup) for picking objects. The robot system (22) receives pick commands and target coordinates from the control unit (20).
- **Bin (24) and Objects in Bin (26):** The container holding the objects to be picked and the objects themselves, which are the target of the depth estimation process.

5. Operation

[0007] The operation of the system proceeds through several phases: pattern projection, image acquisition, depth calculation, disagreement testing, data validation, and pick decision.

[0008] At **system startup**, the control unit (20) initializes the projectors (10, 12) and the imaging sensor (14). Default projection patterns and exposure settings are loaded. The depth processor (16) and validation unit (18) are prepared for data processing.

[0009] During **steady running**, the system continuously cycles through the following steps to acquire and process depth data:

[0010]

- **Pattern Projection and Image Acquisition:**
- The control unit (20) commands Projector 1 (10) to illuminate the scene with its structured light pattern.
- Simultaneously or immediately afterward, the imaging sensor (14) captures an image, denoted as Image A.
- The control unit (20) then commands Projector 2 (12) to illuminate the scene with its structured light pattern.
- The imaging sensor (14) captures a second image, denoted as Image B.
- To account for ambient light, a reference image without any projected pattern may also be captured and subtracted from Image A and Image B.

[0011]

- **Initial Depth Calculation:**
- Image A is sent to the depth processor (16), which calculates a first depth map, Depth Map 1, based on the deformation of Projector 1's (10) pattern observed by the imaging sensor (14).
- Image B is sent to the depth processor (16), which calculates a second depth map, Depth Map 2, based on the deformation of Projector 2's (12) pattern.
- Each depth map contains a set of 3D points (x, y, z) representing the surface of the objects (26) in the bin (24).

[0012]

- **Disagreement Test (Validation):**
- Depth Map 1 and Depth Map 2 are passed to the validation unit (18).

- For each pixel location (u, v) in the imaging sensor's (14) field of view where both Depth Map 1 and Depth Map 2 have a depth reading, the validation unit (18) performs a disagreement test.
- The test compares the depth value `z1` from Depth Map 1 and `z2` from Depth Map 2 for the same projected point (or corresponding scene point).
- If the absolute difference `|z1 - z2|` exceeds a predefined spatial tolerance `T_depth`, the point is considered invalid and is rejected. This tolerance `T_depth` can be set, for example, to a value corresponding to the expected noise or resolution of the depth system, such as 0.5 mm to 2.0 mm for close-range bin picking.
- Points where `|z1 - z2|` is less than or equal to `T_depth` are considered valid and are retained, forming a validated point cloud. Specular reflections tend to generate depth readings that are highly dependent on the projector's angle, causing a significant discrepancy between `z1` and `z2` at the same apparent scene location, thus leading to their rejection.

[0013]

- **Density Check and Pick Decision:**
- The control unit (20) receives the validated point cloud from the validation unit (18).
- It then evaluates the density of valid depth points within regions of interest, typically around potential pick candidates.
- A minimum density threshold, for instance, a ratio of validated points to total points in a local region, or a minimum number of points, is required before a pick attempt is authorized. For example, a candidate pick region may require 75% of its potential points to be validated.
- If the density of valid points is sufficient, the control unit (20) identifies suitable grasp points and sends a pick command, including coordinates and orientation, to the robot system (22).

[0014] **On fault or interruption**, such as insufficient valid depth points:

[0015]

- **Exposure Fallback:** If the density of valid points falls below the required threshold, the control unit (20) initiates a fallback procedure. It adjusts the exposure settings of the imaging sensor (14). This may involve increasing or decreasing the exposure time, or altering the gain, to capture images under different lighting conditions. The system then repeats the pattern projection, image acquisition, depth calculation, and disagreement test steps with the new exposure settings. This helps to compensate for situations where the initial exposure might have been suboptimal for certain surface types, potentially improving the visibility of the structured light patterns on non-specular or diffusely reflective parts of the objects.
- **Repositioning/Re-illumination:** In cases where exposure adjustment alone is insufficient, the control unit (20) may trigger other strategies, such as slight repositioning of the robot end effector to change the viewing angle, or adjusting the intensity of the projectors (10, 12).
- **Error Reporting:** If, after multiple attempts with adjusted parameters, the system cannot acquire a sufficient density of valid points, it can report an error condition to a higher-level system or operator, indicating an unpickable state for the current bin configuration.

6. Example Embodiments

[0016] **Embodiment 1** This embodiment utilizes two structured light projectors (10, 12) positioned symmetrically with respect to the imaging sensor (14) and oriented to project patterns from different horizontal angles. For example, if the imaging sensor (14) is centered above the bin (24) looking downwards, Projector 1 (10) might be positioned to the left at an angle of 30 degrees from the sensor's optical axis, and Projector 2 (12) to the right at a corresponding 30 degrees. The system operates by acquiring images sequentially for each projector's pattern. The disagreement test in the validation unit (18) considers points validated if their calculated depth difference is within 1.5 mm. The minimum density requirement for a pick is 80% validated points within a 10 mm radius of the proposed grasp point. If this is not met, the control unit (20) increases the imaging sensor's (14) exposure time by 20% and re-attempts the measurement cycle.

[0017] **Embodiment 2** In this configuration, the two projectors (10, 12) are positioned such that one projects from a high angle (e.g., 45 degrees from vertical) and the other from a low angle (e.g., 20 degrees from vertical), both offset in the same horizontal direction relative to the imaging sensor (14). This angular separation in the vertical plane provides differing perspectives on potential specular reflections. The projectors (10, 12) use different structured light patterns; for instance, Projector 1 (10) projects a horizontal line pattern, and Projector 2 (12) projects a vertical line pattern. The depth processor (16) is capable of processing both pattern types. The validation unit (18) uses a dynamic tolerance for the disagreement test, where `T_depth` varies with the estimated distance from the sensor, allowing a larger tolerance for farther points (e.g., `T_depth = 0.01 * Z`, where Z is the depth). The control unit (20) manages exposure by attempting a short exposure (e.g., 5 ms) first, and if insufficient valid points are obtained, it switches to a longer exposure (e.g., 25 ms) for a second attempt.

[0018] **Embodiment 3** This embodiment employs a single physical structured light projector (10) with two distinct projection modes, effectively acting as two virtual projectors (10, 12) from different angles. This can be achieved using a programmable spatial light modulator (e.g., a Digital Micromirror Device) to rapidly switch between two distinct pattern orientations or projection directions. For instance, the projector (10) could project a pattern from a slightly shifted virtual position by steering the light. The imaging sensor (14) captures images for each virtual projection angle. The disagreement test in the validation unit (18) is augmented to also consider the consistency of surface normals derived from the two depth maps; if the normals diverge significantly (e.g., by more than 10 degrees), the point is also rejected. The minimum density requirement is applied across the entire visible surface of an object, not just potential pick points, to ensure a comprehensive object model. The fallback strategy involves increasing the projector's (10) light intensity by a factor of 1.5 before adjusting exposure, followed by a second attempt with modified exposure if still insufficient.

7. Alternative Configurations

[0019] The described system can be configured with various alternatives. Instead of two discrete projectors (10, 12), a single projector with a movable mirror or a programmable spatial light modulator could effectively generate patterns from two distinct viewpoints or angles. The structured light patterns themselves can vary; while lines or grids are common, patterns such as pseudo-random dots or coded light patterns (e.g., Gray code sequences) could also be employed. If coded light patterns are used, the system would project multiple frames per "projector" to encode depth information, and the disagreement test would apply to the final reconstructed depth maps. The imaging sensor (14) could be a monochrome or color camera, or even a specialized sensor like a time-of-flight camera for initial coarse depth estimation, which is then refined by the structured light technique. The control unit (20) could integrate machine learning algorithms to predict optimal exposure settings or pattern types based on real-time analysis of object surface characteristics. The disagreement test logic could also incorporate confidence metrics from the depth calculation, weighting points with higher confidence more heavily.

8. Parameter Variations

Parameter	Realistic Range	Effect of Moving Within Range
Projector Angular Separation (α, β)	20 to 60 degrees	Larger angles increase baseline for depth calculation and enhance specular rejection but may create more shadows. Smaller angles reduce shadows but decrease the effectiveness of specular rejection.
Working Distance (D)	0.3 to 2.0 meters	Shorter distances improve depth resolution and accuracy. Longer distances increase field of view but reduce resolution and increase susceptibility to environmental light.
Spatial Resolution of Sensor	0.1 to 1.0 mm at working distance	Higher resolution captures finer details and improves accuracy for small features. Lower resolution reduces computational load but may miss critical object features.
Depth Disagreement Tolerance (T_depth)	0.5 to 5.0 mm	Tighter tolerance rejects more points, including some valid but noisy ones, increasing robustness against specularity but potentially reducing data density. Looser tolerance retains more points but increases false positives from specularities.
Minimum Valid Point Density	60% to 95% of region	Higher density threshold ensures more reliable pick points but may lead to more frequent exposure fallback cycles. Lower threshold allows picks with less data but increases risk of failure.
Exposure Time	1 to 50 milliseconds	Longer exposure times capture more light, improving signal-to-noise ratio for dark surfaces, but increase motion blur sensitivity. Shorter exposures reduce blur but may underexpose dark areas.
Projector Light Intensity	100 to 1000 lumens	Higher intensity improves visibility of patterns on dark or distant objects but can cause saturation on highly reflective surfaces. Lower intensity reduces saturation but may not be visible enough.

9. Additional Implementations

[0021] The principles of this dual-angle structured light system with specular rejection and exposure fallback can be extended to other 3D sensing applications beyond bin picking. For example, it could be implemented in quality inspection systems to accurately measure the geometry of metallic or polished parts, where conventional structured light systems struggle. Another implementation could involve robotic assembly, where precise localization of mating parts with varied surface finishes is crucial. The system could also be adapted for autonomous navigation of vehicles or robots in environments with highly reflective surfaces, providing more reliable obstacle detection. Furthermore, the disagreement test methodology can be generalized to incorporate more than two projectors, each at a unique angle, for even greater redundancy and robustness against complex specular behaviors or inter-reflections. The fallback mechanism could include a dynamic pattern selection, where the control unit (20) chooses from a library of structured light patterns (e.g., different frequencies of lines, or random dots) based on the initial assessment of surface properties, before resorting to exposure changes. The system could also incorporate a diffuse illumination phase (without structured light) to capture color texture information, which can then be mapped onto the validated 3D point cloud for improved object recognition and grasp planning.

10. Figures

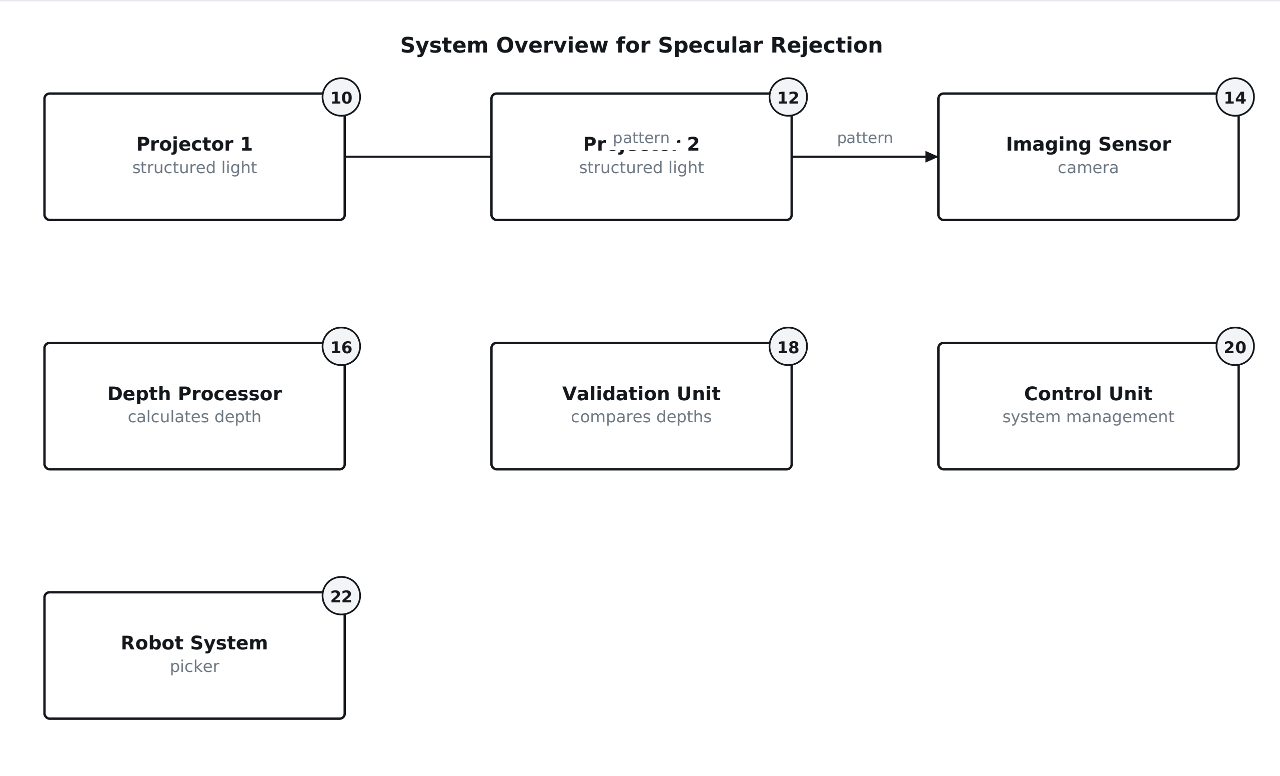

Interactions between hardware and software modules.

FIG. 1 A high-level block diagram showing the flow of data and control within the depth estimation system. PNG

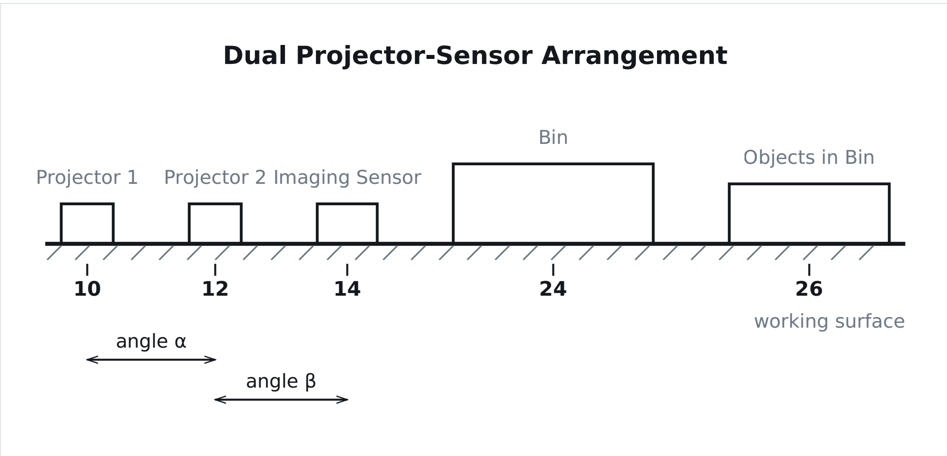

Angles are relative to the sensor-object axis.

FIG. 2 An elevation view illustrating the geometric setup of the two structured light projectors and the imaging sensor relative to a bin of objects. PNG

OpenTechnical Publication OT-2026-000000047. Published September 15, 2026 at 11:48:03 UTC. https://opentechnical.org/p/OT-2026-000000047