Algorithms and Systems for Low Power Time-of-Flight Imaging

James Noraky

April 17, 2020

Thesis Advisor: **Prof. Vivienne Sze**

Thesis Committee: **Prof. Berthold K.P. Horn**

Dr. Charles Mathy

Depth Sensors Enable Many Emerging Applications

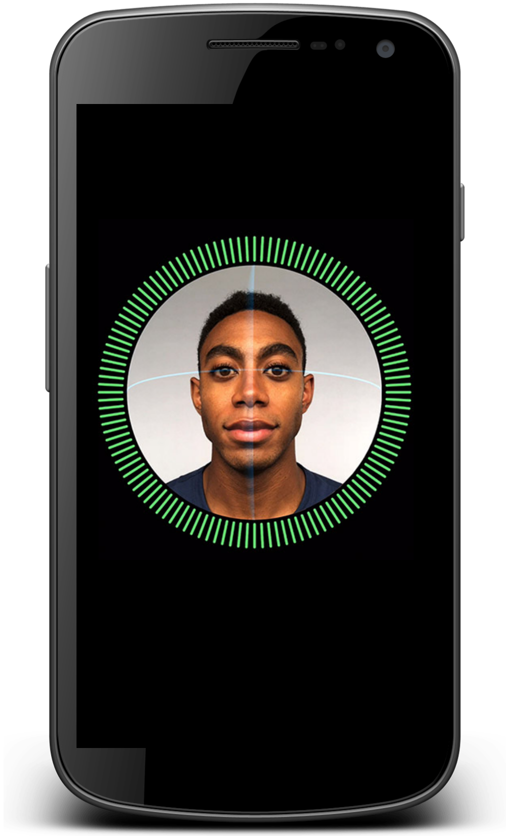

Depth information enables **safe navigation** and **interactive applications**

Time-of-Flight Cameras Are Appealing Depth Sensors

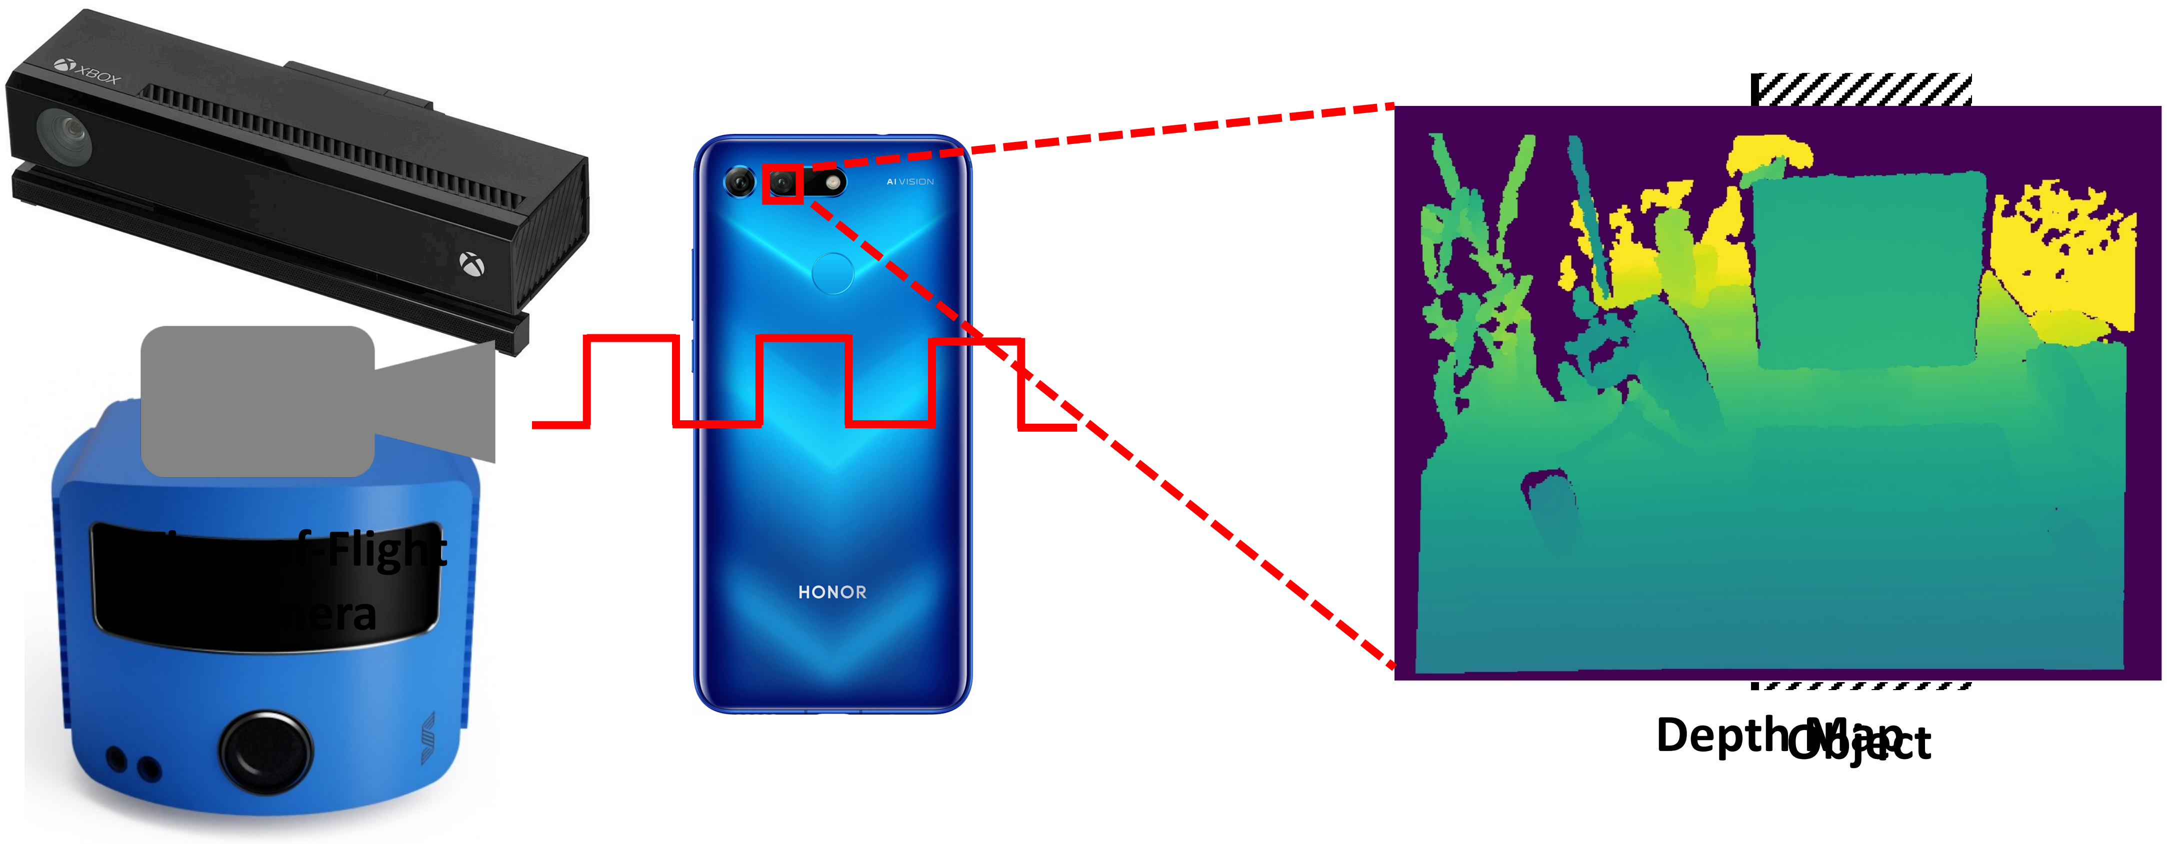

Obtain depth by emitting light and measuring its round trip travel time

No moving parts, compact, obtain dense depth maps with minimal latency

Time-of-Flight Cameras Are Power Hungry

- **Active Sensor**: For ranges up to 8 m, time-of-flight (ToF) cameras consume up to 20 W
- **Reduced Battery Life**: Especially for applications that need continuous depth
- **Increased Heat Dissipation**: Affects calibration and forces the addition of bulky heat sinks

Thesis Goals

- Lower the sensor power of ToF cameras
- Obtain accurate and dense depth maps
- Minimize the latency of our approaches on low power processors

Reduce the Usage of the ToF Camera

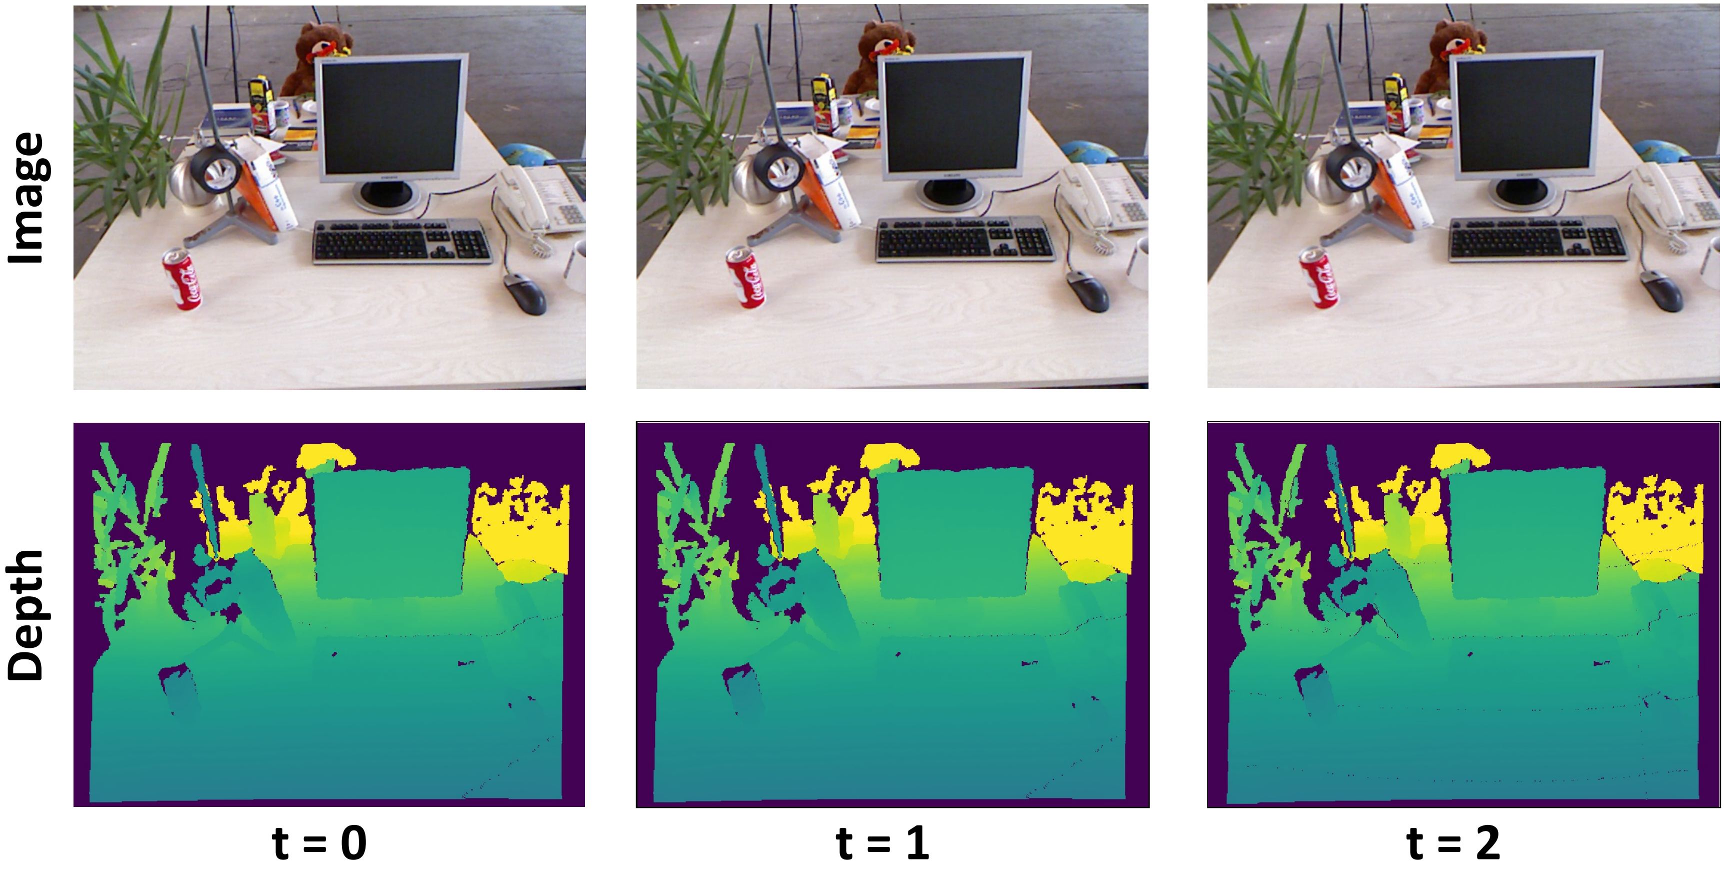

Estimate depth maps using consecutive and concurrently collected RGB images

Reduce the Light the ToF Camera Emits

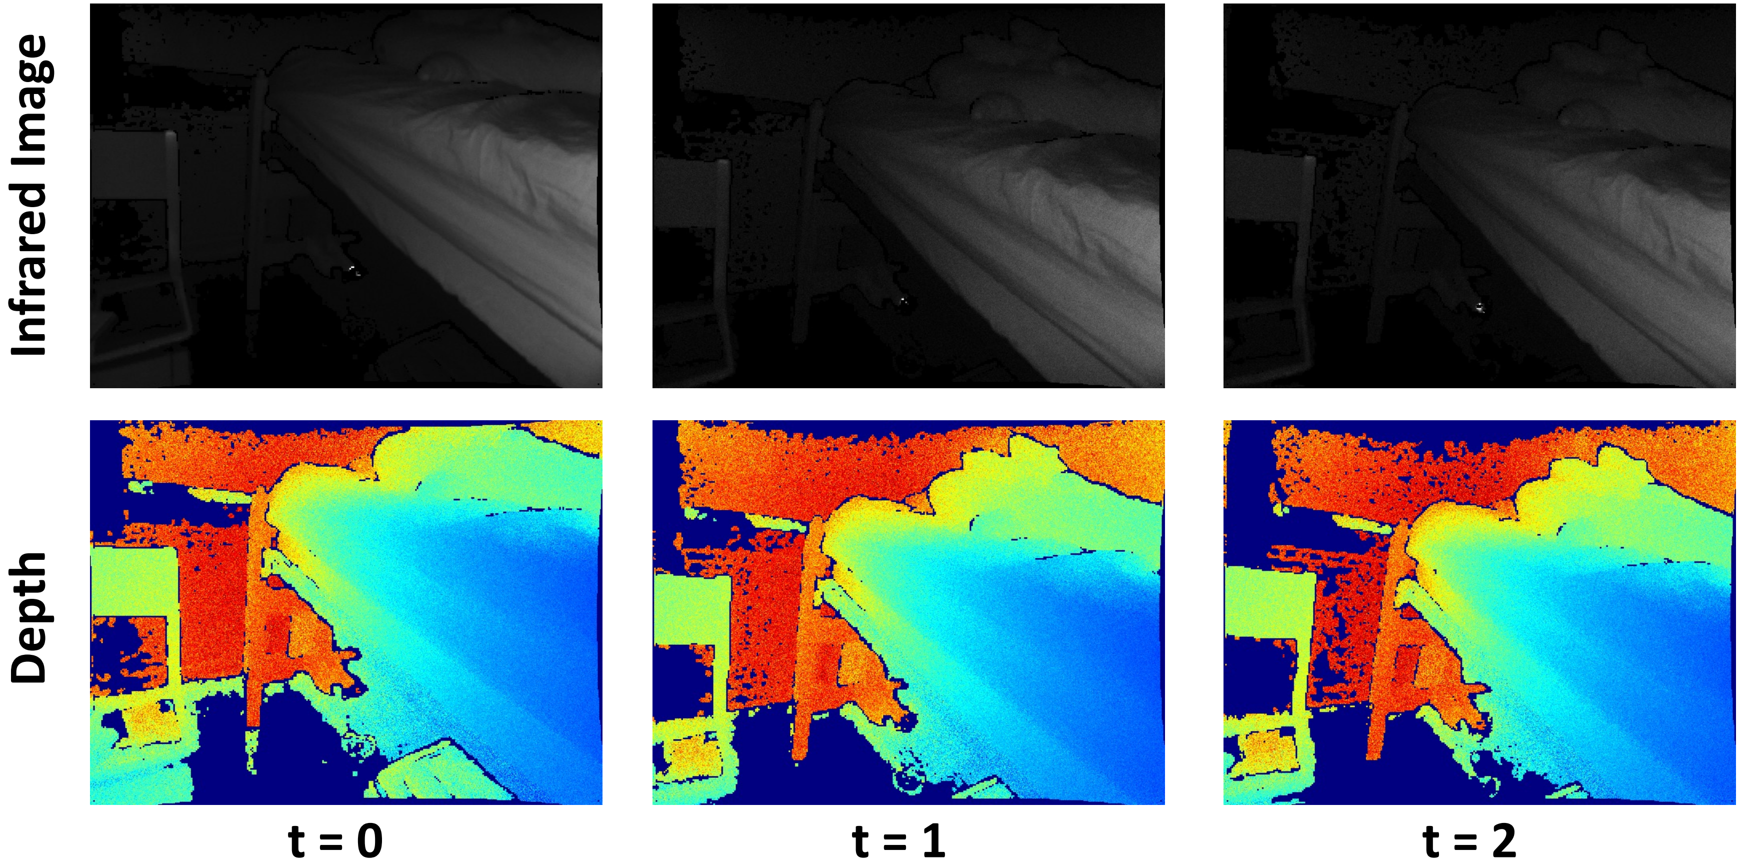

Denoise the low power depth maps by combining depth maps across frames

Balance Sensor Power, Accuracy, and Latency

Strictly Computation

Strictly ToF Camera

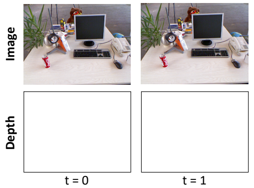

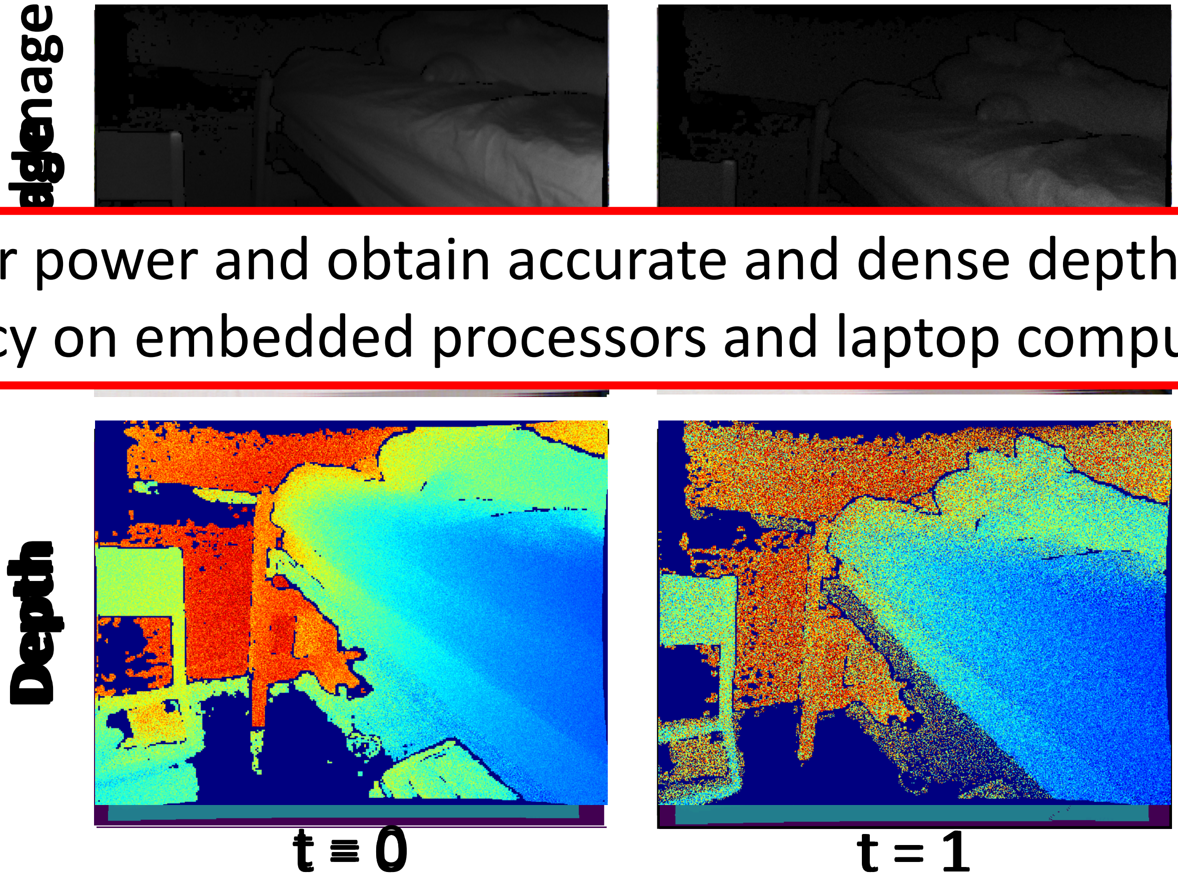

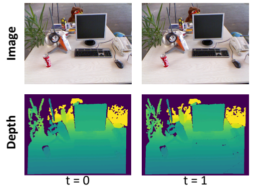

Low S

Low

High Latency

Lower the sensor power and obtain accurate and dense depth maps with low latency on embedded processors and laptop computers

Power

acy

Low Latency

The Metrics We Use to Evaluate Our Approaches

- **Sensor Power**
  - Duty Cycle: percentage of frames where the ToF camera is used
  - Normalized Power: fraction of light emitted compared to a regular depth map

- **Accuracy**
  - Mean Relative Error (MRE) $= \frac{100}{N} \sum_{i=1}^{N} \frac{|Z_i - \hat{Z}_i|}{Z_i}$
  - N is the total number of pixels, Z_i is the ground truth depth, $\hat{Z}_i$ is the estimated depth

- **Latency**
  - Quantify the estimation frame rate (FPS) on a low power embedded processor and laptop computer

Estimating Accurate and Dense Depth Maps Is Hard

- **Challenge**: Problem is underdetermined and requires dense computation
- Assume that the scenes contain **rigid motions** and use them to estimate/denoise depth maps
- Rigid motion can be efficiently estimated using **sparse operations** and **linear least squares**

Publications That Went Into Thesis

- **Conferences**
  - J. Noraky, V. Sze, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
  - J. Noraky, V. Sze, "Depth Estimation of Non-Rigid Objects for Time-of-Flight Imaging," ICIP, 2018.
  - J. Noraky, C. Mathy, A. Cheng, V. Sze, "Low Power Adaptive Time-of-Flight Imaging for Multiple Rigid Objects," ICIP, 2019.

- **Journal Publications and Preprints**
  - J. Noraky, V. Sze, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.
  - J. Noraky, V. Sze, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.
  - J. Noraky, V. Sze, “Low Power Depth Map Denoising for Mobile Time-of-Flight Cameras,” In Preparation, 2020.

Publications That Went Into Thesis

- **Conferences**
  - J. Noraky, V. Sze, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
  - J. Noraky, V. Sze, "Depth Estimation of Non-Rigid Objects for Time-of-Flight Imaging," ICIP, 2018.
  - J. Noraky, C. Mathy, A. Cheng, V. Sze, "Low Power Adaptive Time-of-Flight Imaging for Multiple Rigid Objects," ICIP, 2019.

- **Journal Publications and Preprints**
  - J. Noraky, V. Sze, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.
  - J. Noraky, V. Sze, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.
  - J. Noraky, V. Sze, "Low Power Depth Map Denoising for Mobile Time-of-Flight Cameras," In Preparation, 2020.

Outline

Strictly Computation

Strictly ToF Camera

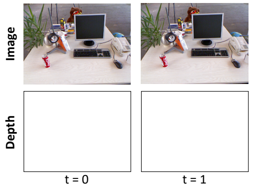

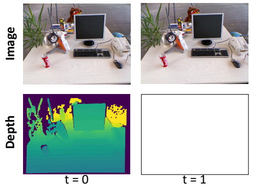

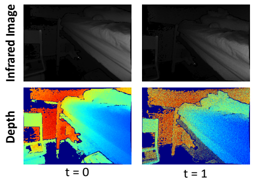

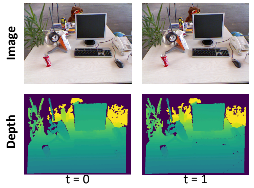

- Reduce the Usage of the ToF Camera
  - Depth Map Estimation for Rigid Scenes
  - Depth Map Estimation for Dynamic Scenes
- Reduce the Light the ToF Camera Emits
  - Adaptive Pulse Control
- Summary of Thesis Contributions

Outline

Strictly Computation

Strictly ToF Camera

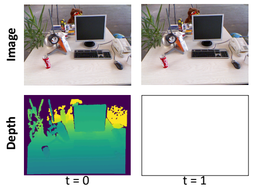

- Reduce the Usage of the ToF Camera
  - **Depth Map Estimation for Rigid Scenes**
  - Depth Map Estimation for Dynamic Scenes
- Reduce the Light the ToF Camera Emits
  - Adaptive Pulse Control
- Summary of Thesis Contributions

Depth Map Estimation for Rigid Scenes

- Estimate the depth map in the current frame
- Between consecutive frames, there is not a lot of motion
- Update the previous depth map using motion cues from the images

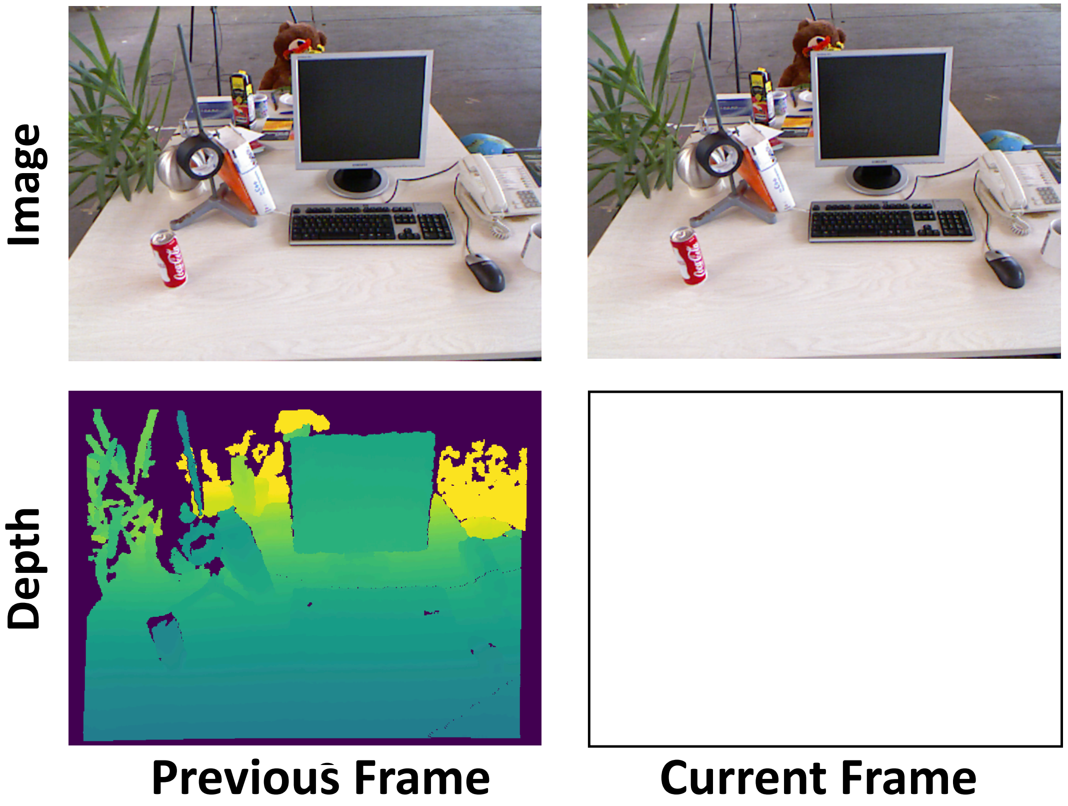

Depth Map Estimation for Rigid Scenes

- Many approaches use the dense optical flow between the images to remap the pixels of the previous depth map

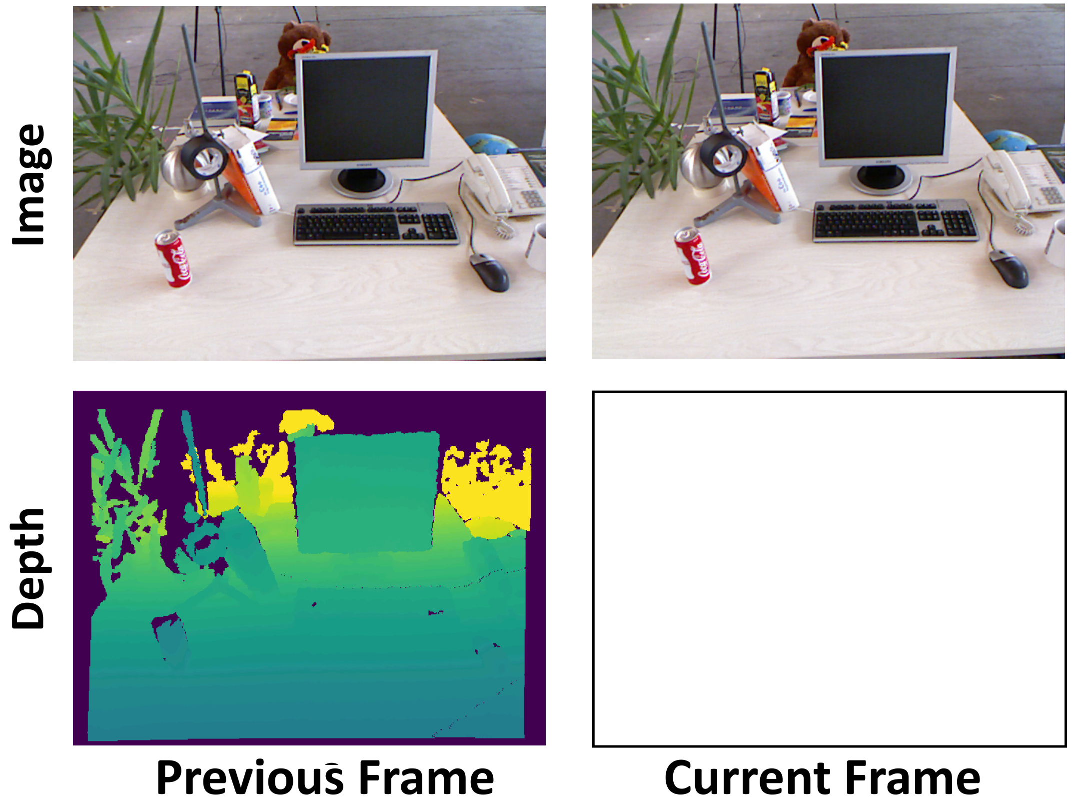

Optical Flow Is the Apparent Pixel-Wise Motion

Previous Frame

Current Frame

Estimating dense and accurate optical flow is computationally expensive

Use the optical flow to move the pixels of the previous depth map

Depth Map Estimation for Rigid Scenes

- Many approaches use the **dense optical flow** between the images to remap the pixels of the previous depth map

- Estimate the **rigid motion using sparse optical flow** and use it to reproject the previous depth map

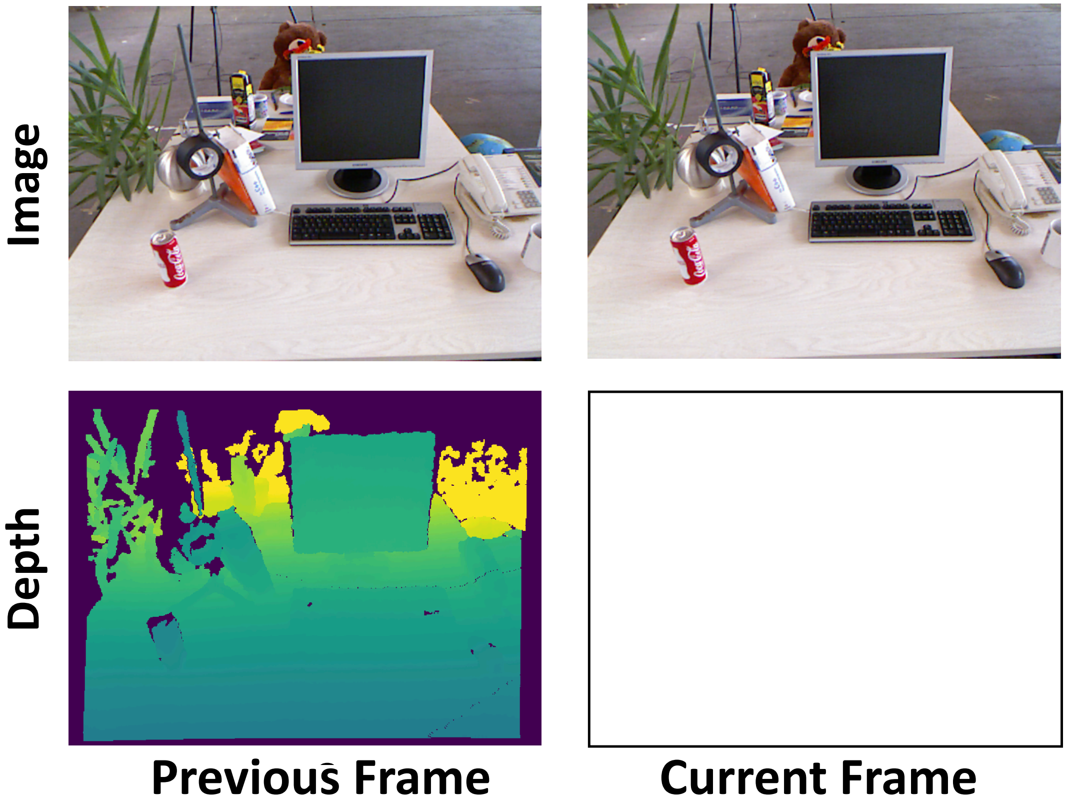

Estimating the Rigid Motion

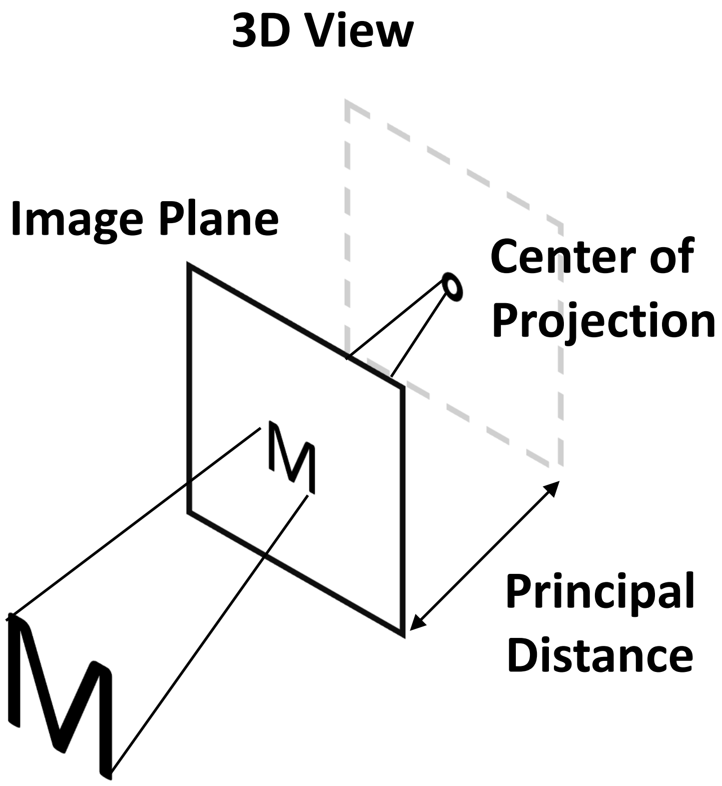

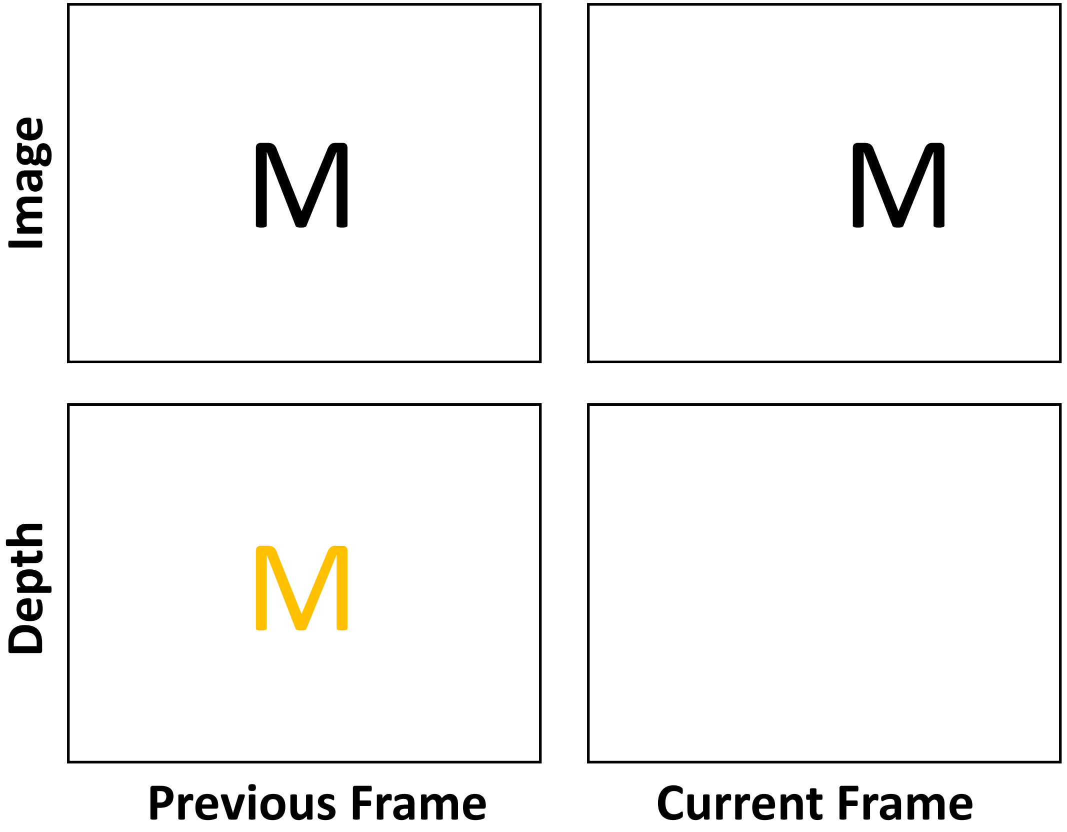

Given depth in the previous frame, we can obtain the 3D location of each pixel

Assume perspective projection

Estimating the Rigid Motion

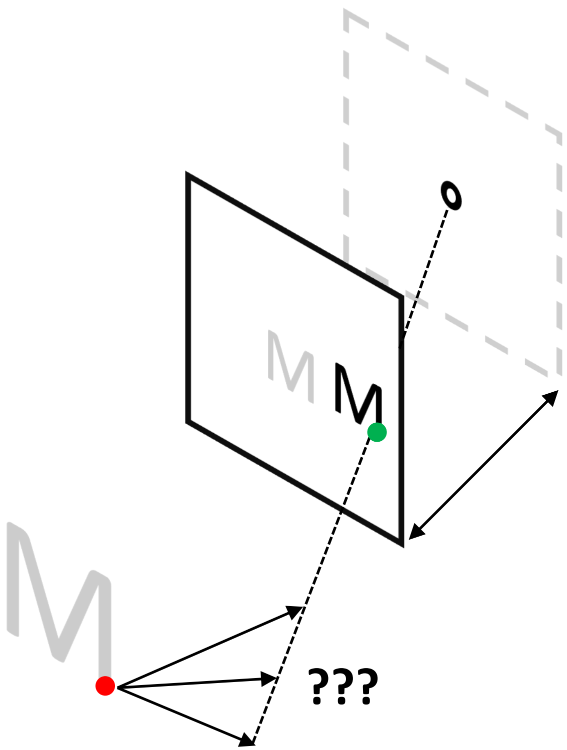

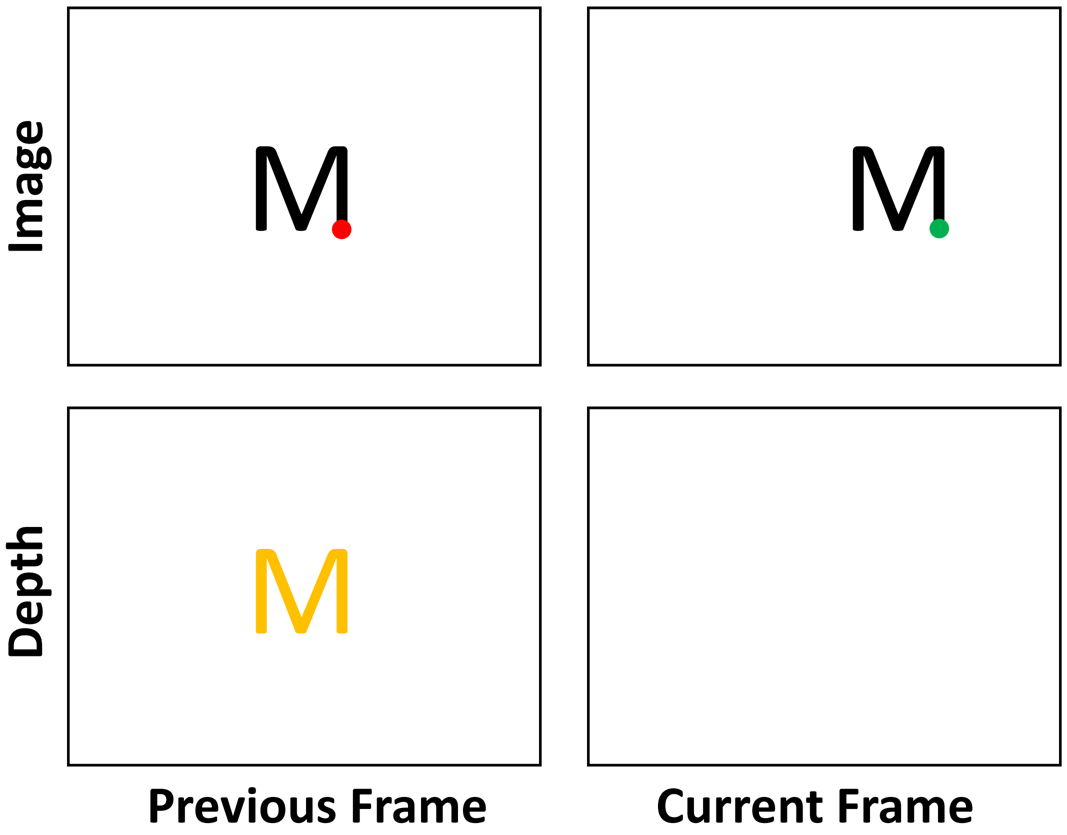

What does the perspective projection tell us?

How does the "M" move in 3D?

Both are ambiguous!

Estimating the Rigid Motion

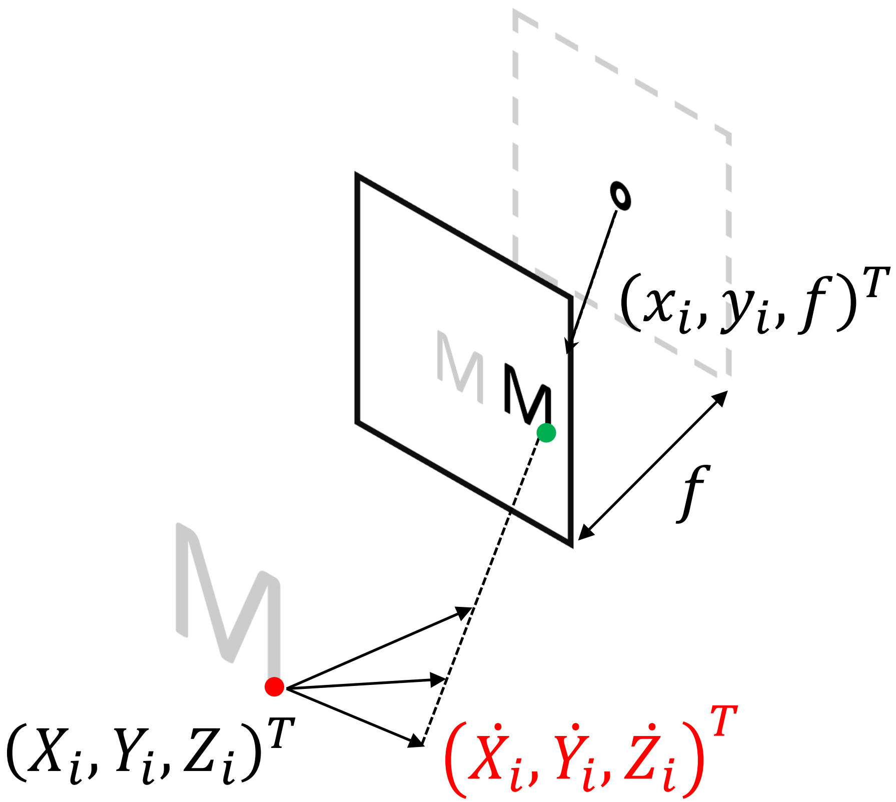

$$\left(\begin{bmatrix} X_i \\ Y_i \\ Z_i \end{bmatrix} + \begin{bmatrix} \dot{X}_i \\ \dot{Y}_i \\ \dot{Z}_i \end{bmatrix}\right) \times \begin{bmatrix} x_i \\ y_i \\ f \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

3D coordinate in the previous depth map

3D displacement due to rigid motion

Ray containing the correspondence in the current frame

2 independent equations with 3 unknowns for each pixel

Use the rigidity assumption

Estimating the Rigid Motion

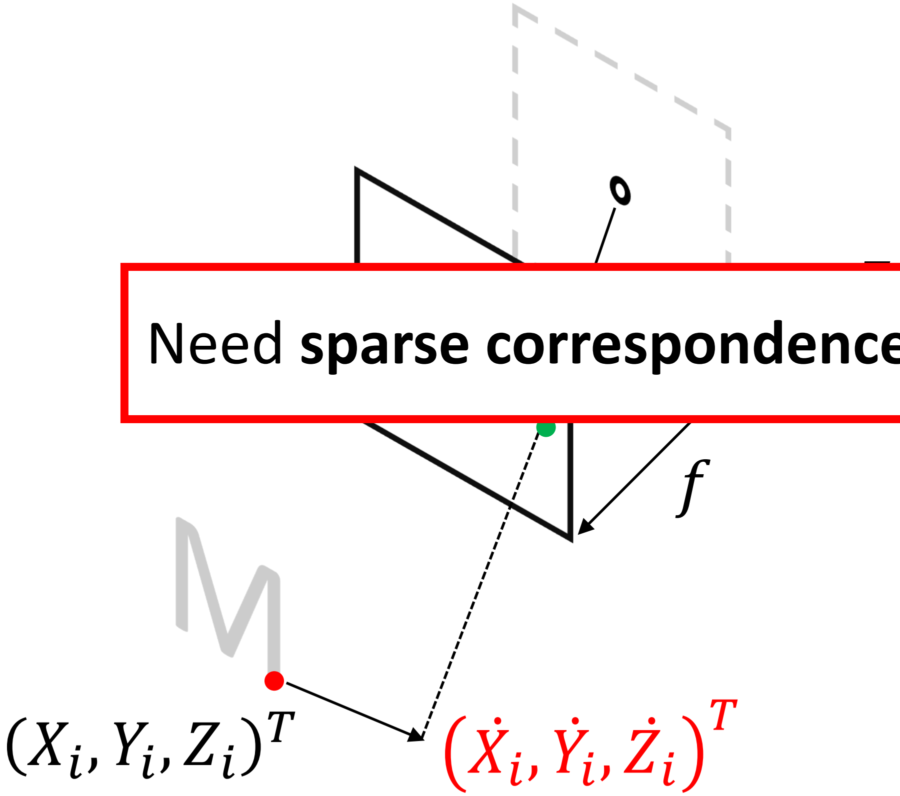

$$\left(\begin{bmatrix} X_i \\ Y_i \end{bmatrix} + \boldsymbol{\omega} \times \begin{bmatrix} X_i \\ Y_i \end{bmatrix} + \boldsymbol{T}\right) \times \begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

Need **sparse correspondences** to estimate rigid motion with **linear least squares**

Velocity Velocity

2 independent equations in 6 unknowns

Depth Map Estimation for Rigid Scenes

- Many approaches use the dense optical flow between the images to **remap** the pixels of the previous depth map

- Estimate the rigid motion using sparse optical flow and use it to **reproject the previous depth map**

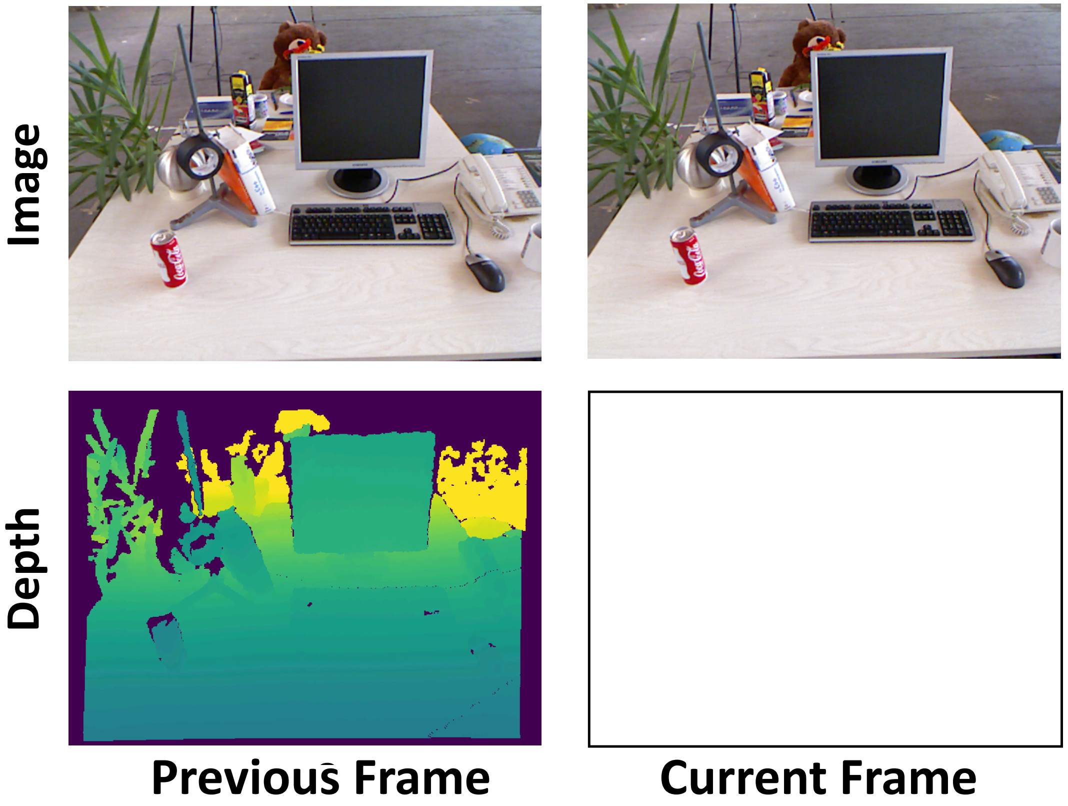

Estimating the Rigid Motion

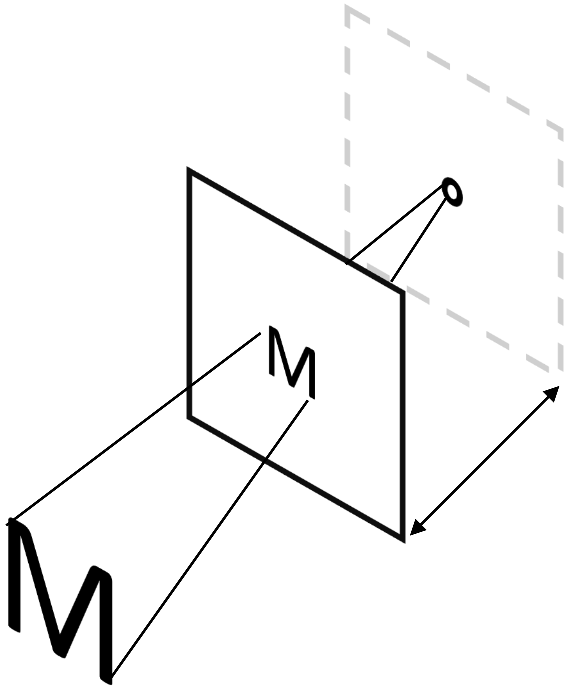

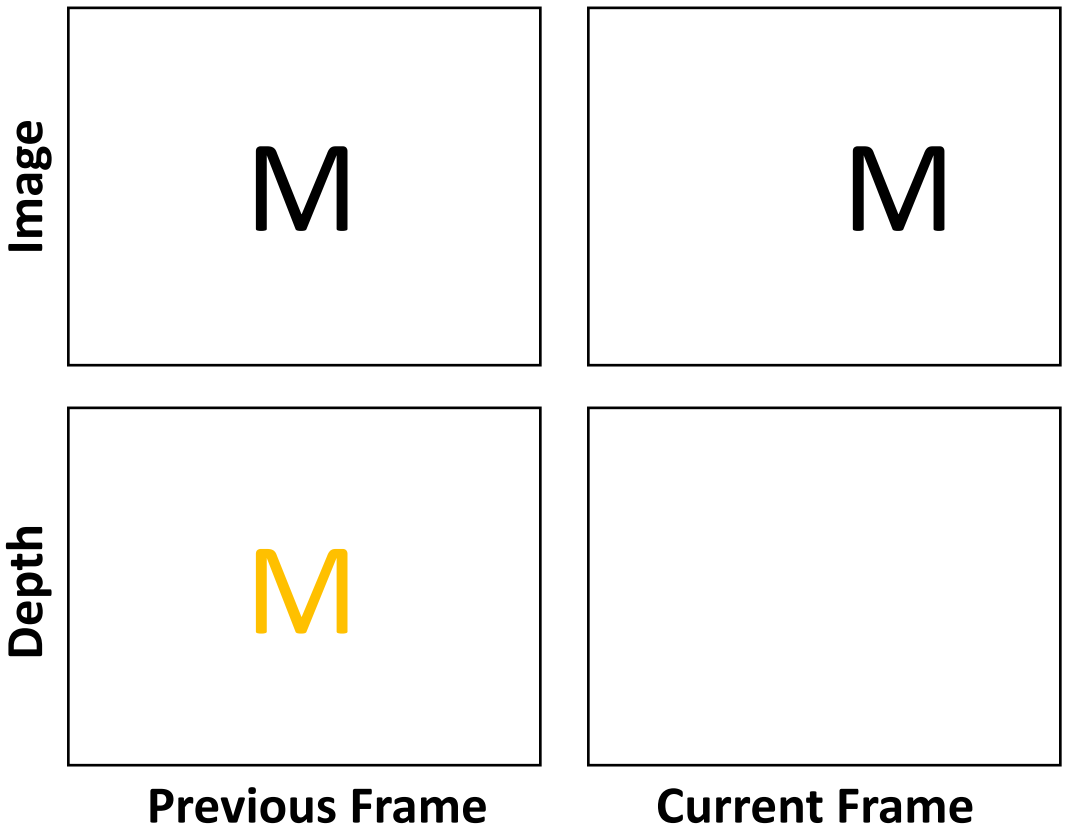

Get the 3D position of each point

Estimating the Rigid Motion

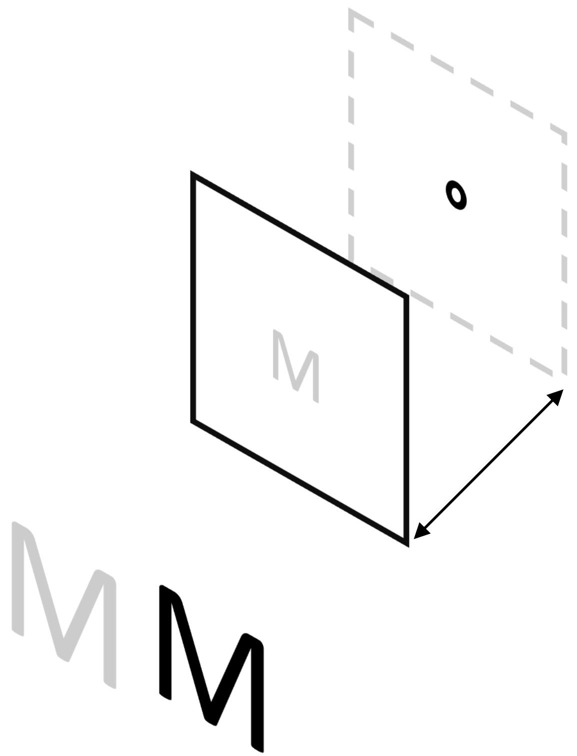

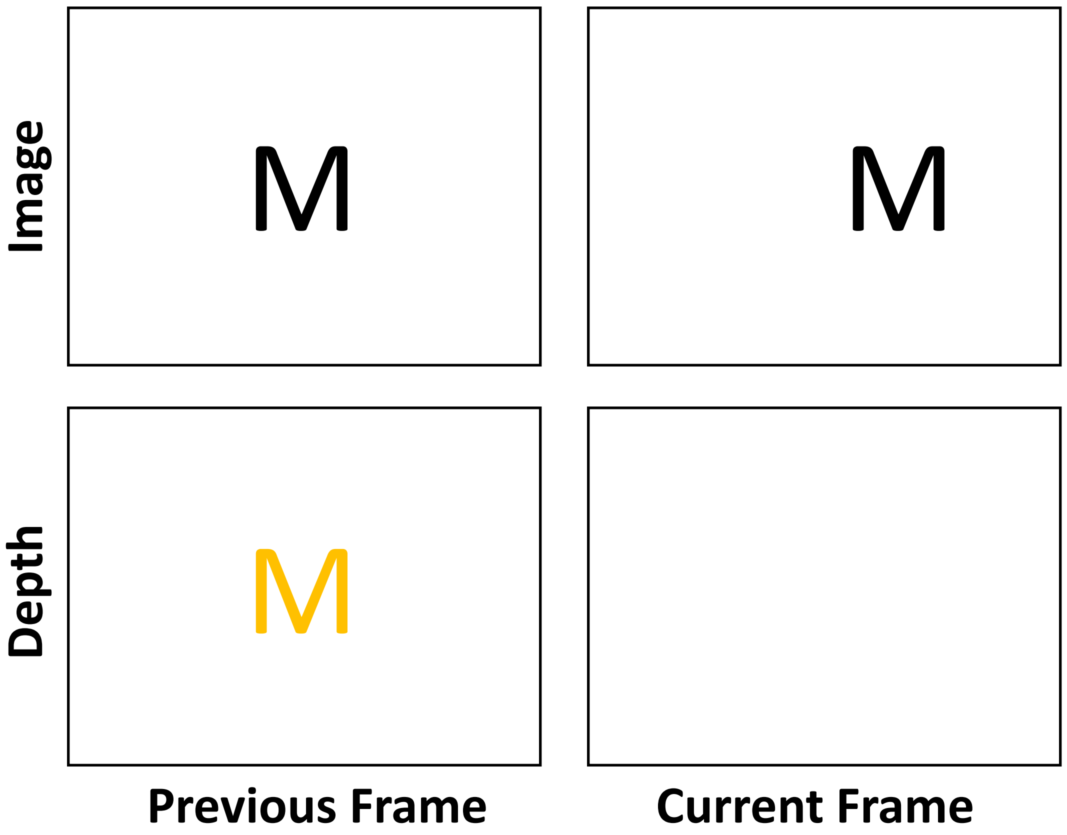

Apply the rigid motion to the 3D points

Estimating the Rigid Motion

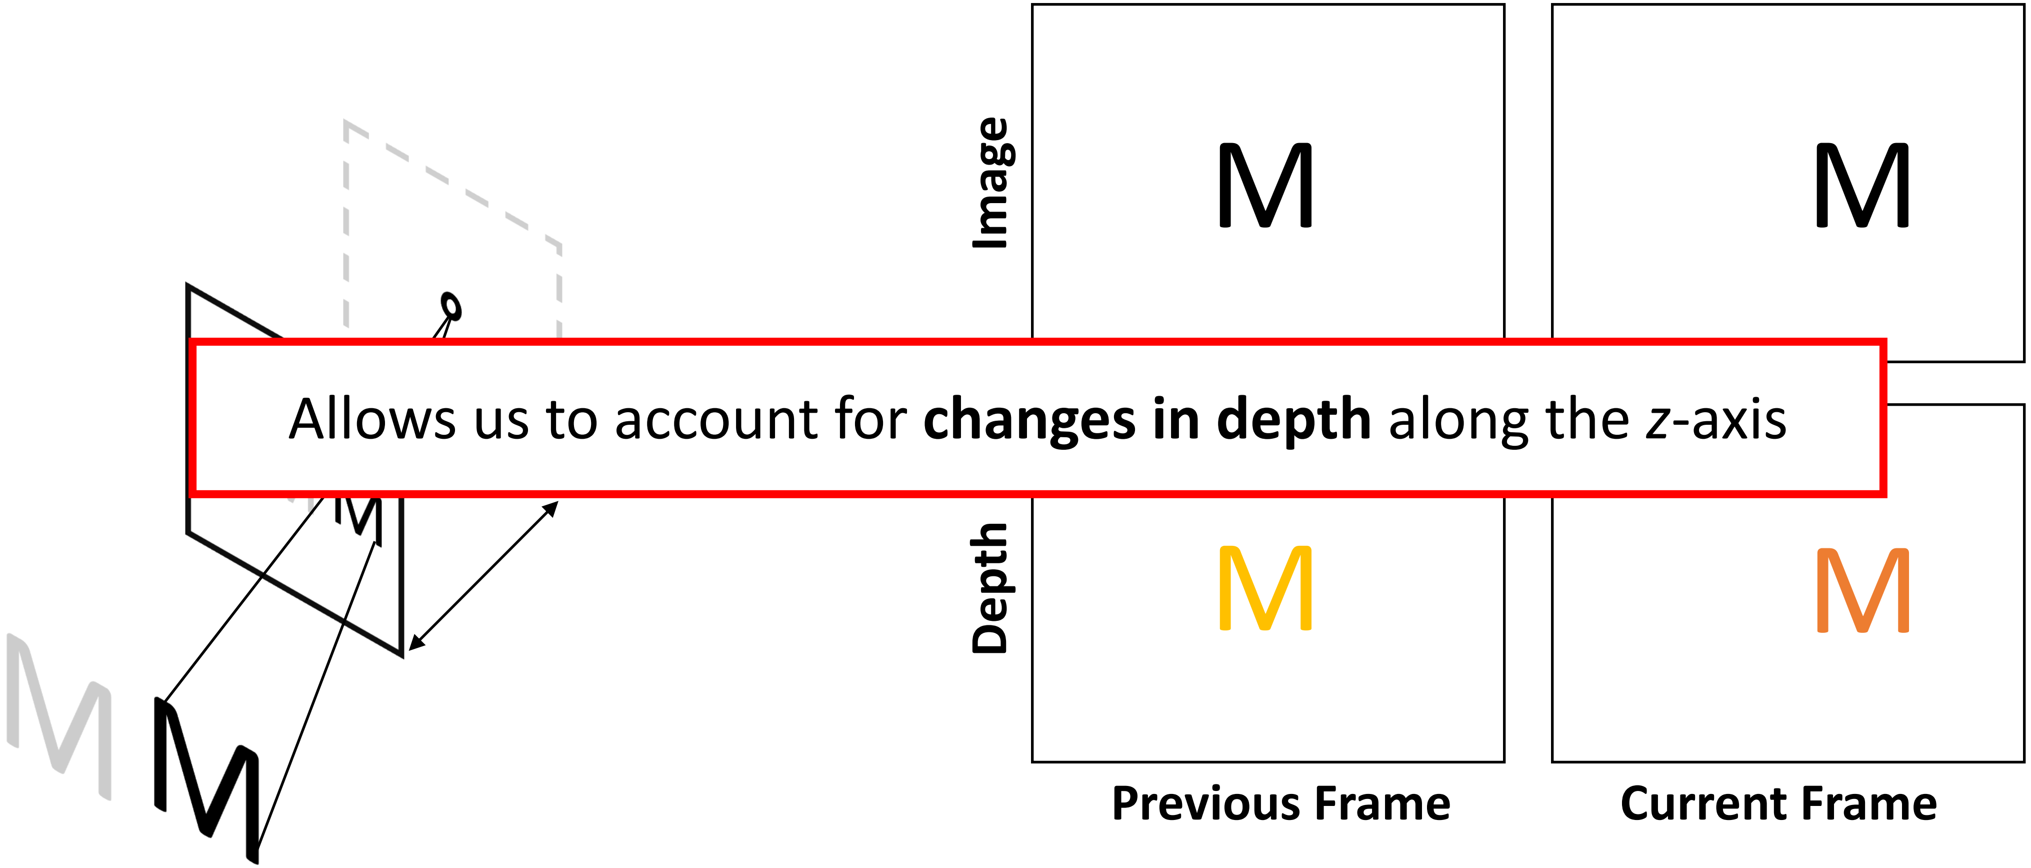

Allows us to account for **changes in depth** along the *z*-axis

Project the **updated depth** to the image

Depth Map Estimation for Rigid Scenes

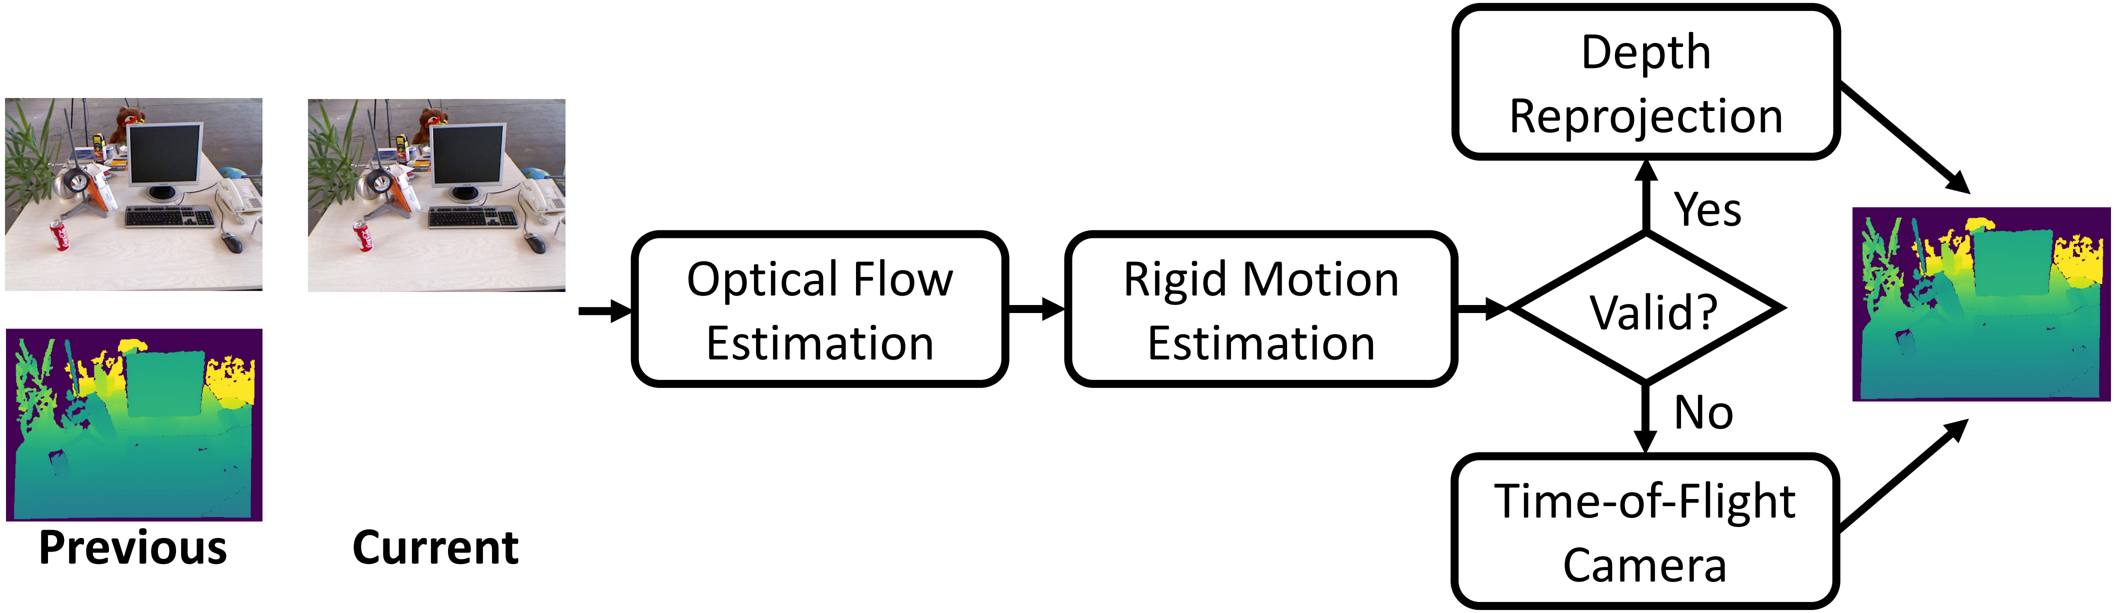

Noraky *et al.*, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
Noraky *et al.*, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.

Depth Map Estimation for Rigid Scenes

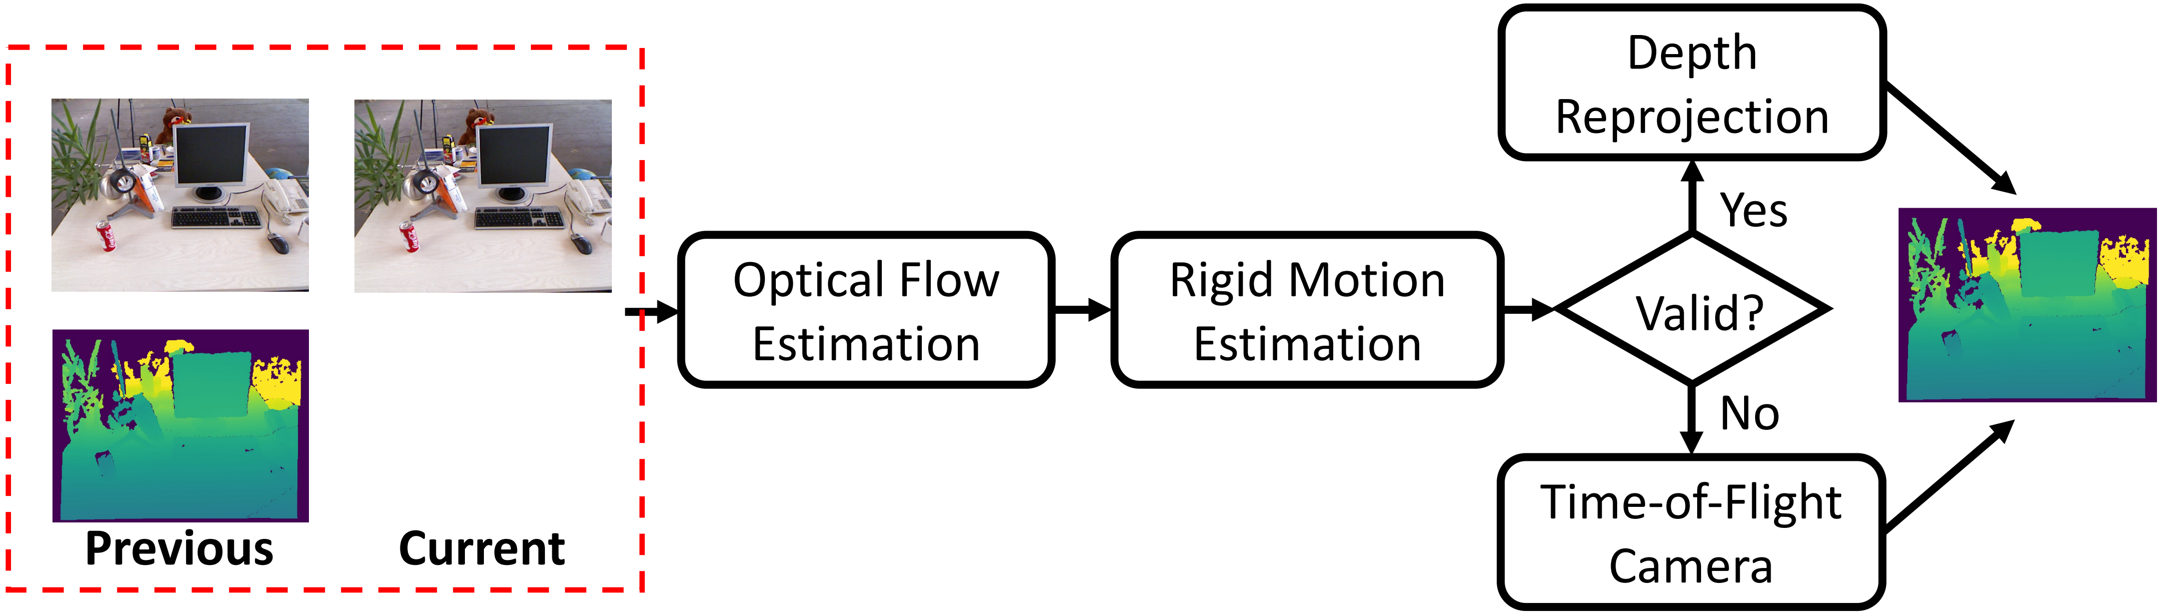

Our inputs are consecutive images and a previous depth map

Noraky *et al.*, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
Noraky *et al.*, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.

Depth Map Estimation for Rigid Scenes

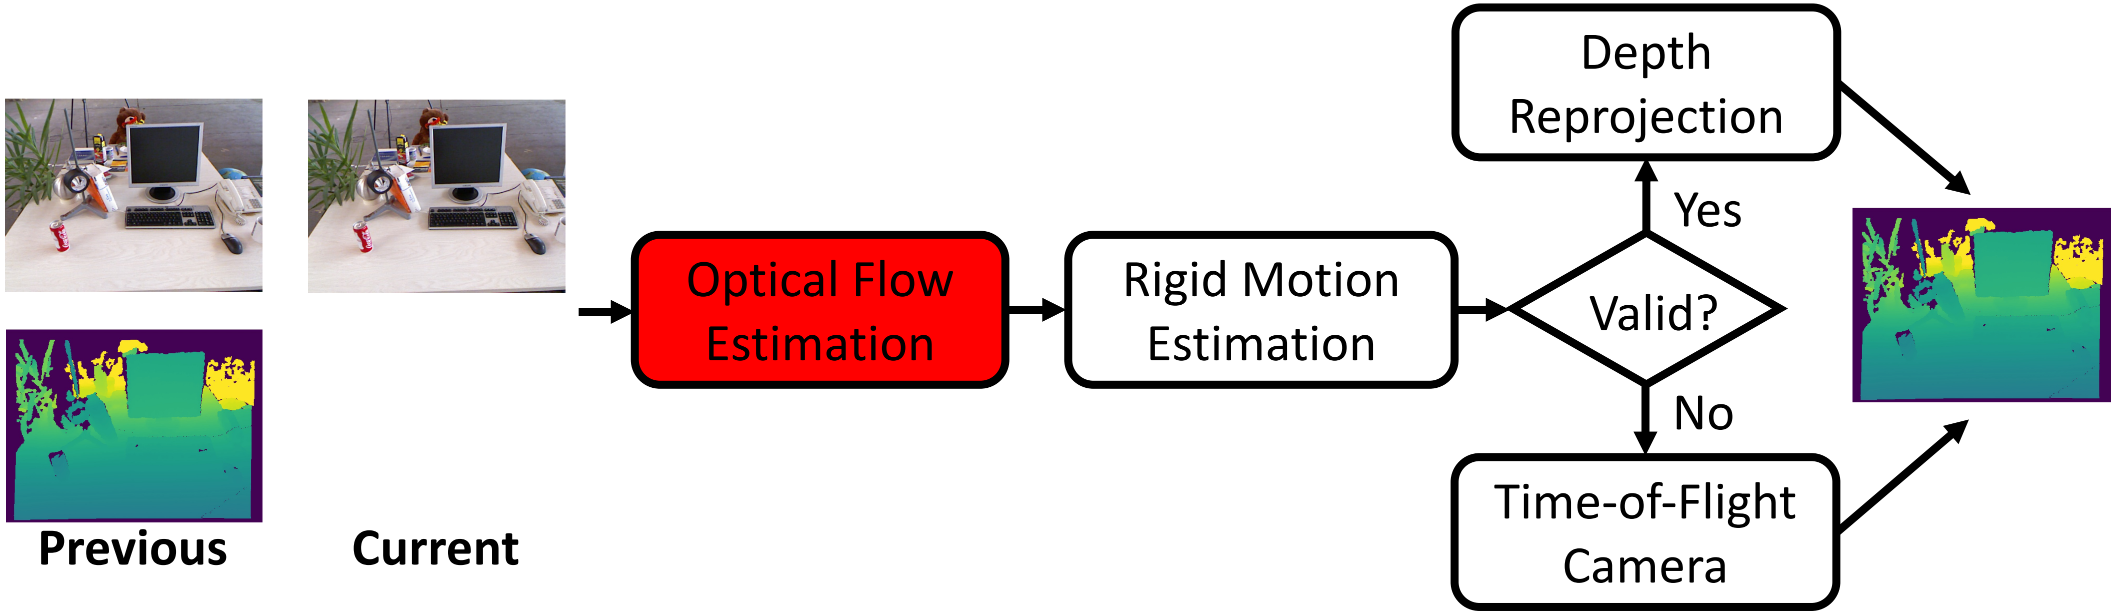

Use **block matching** on the pixels on a **sparse grid**

Noraky *et al.*, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
Noraky *et al.*, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.

Use Efficient Block Matching Heuristic

Previous Frame

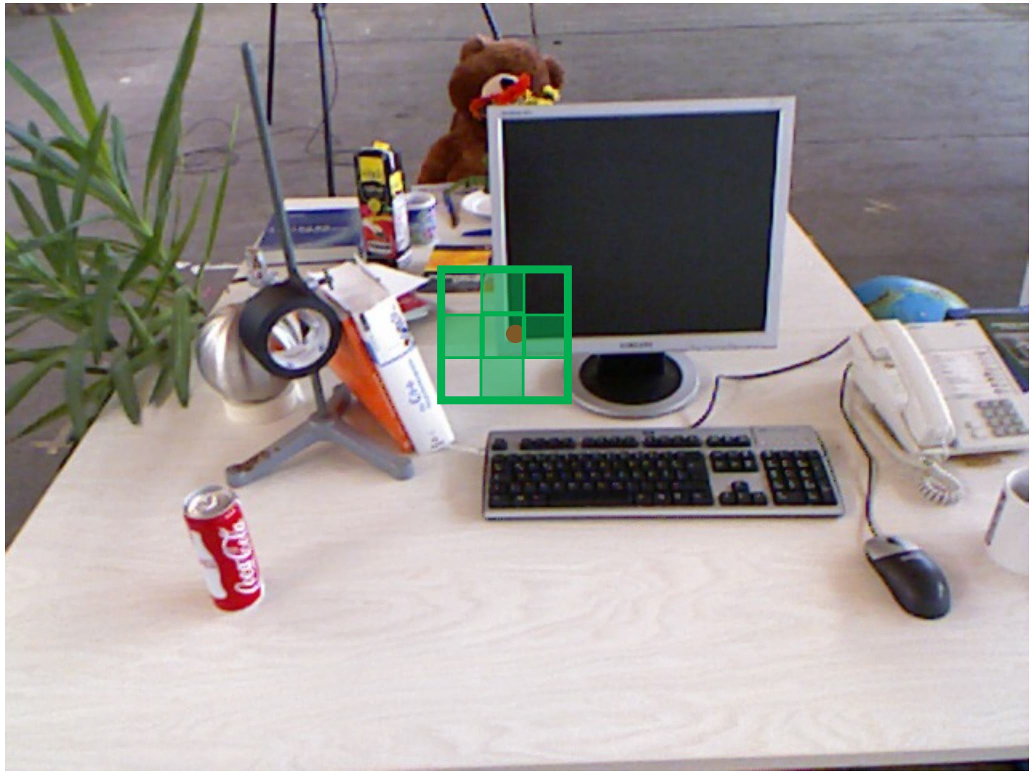

Current Frame

Reduce the number of positions used to determine the correspondence

Depth Map Estimation for Rigid Scenes

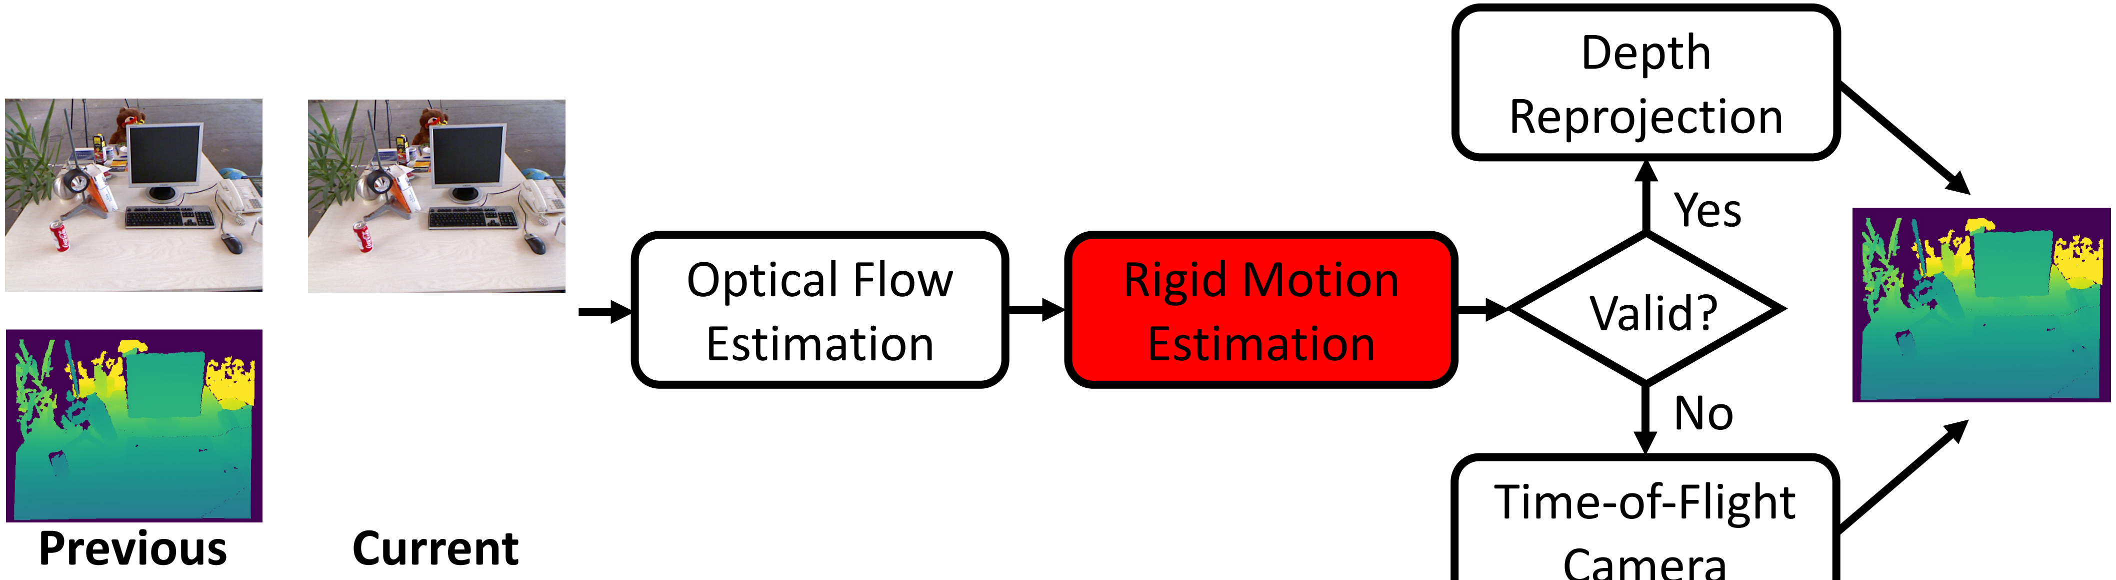

Estimate the rigid motion robustly using RANSAC

Noraky *et al.*, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
Noraky *et al.*, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.

RANSAC Mitigates the Impact of Outliers

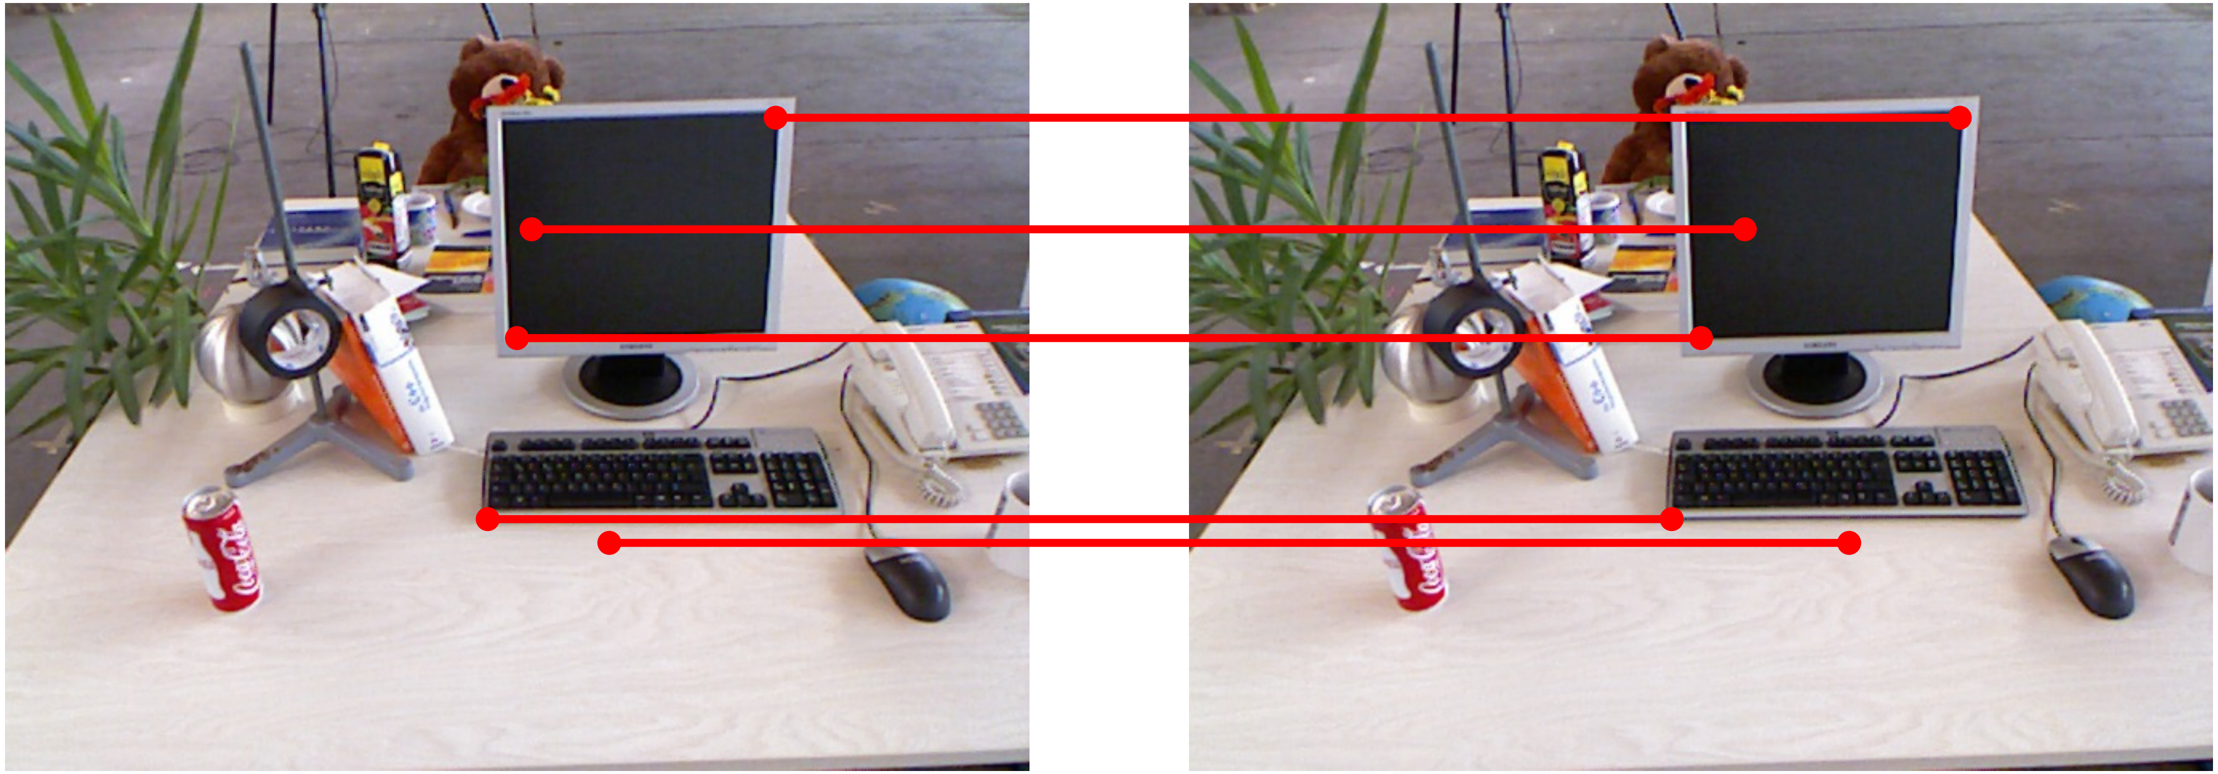

Previous Frame

Current Frame

RANSAC Mitigates the Impact of Outliers

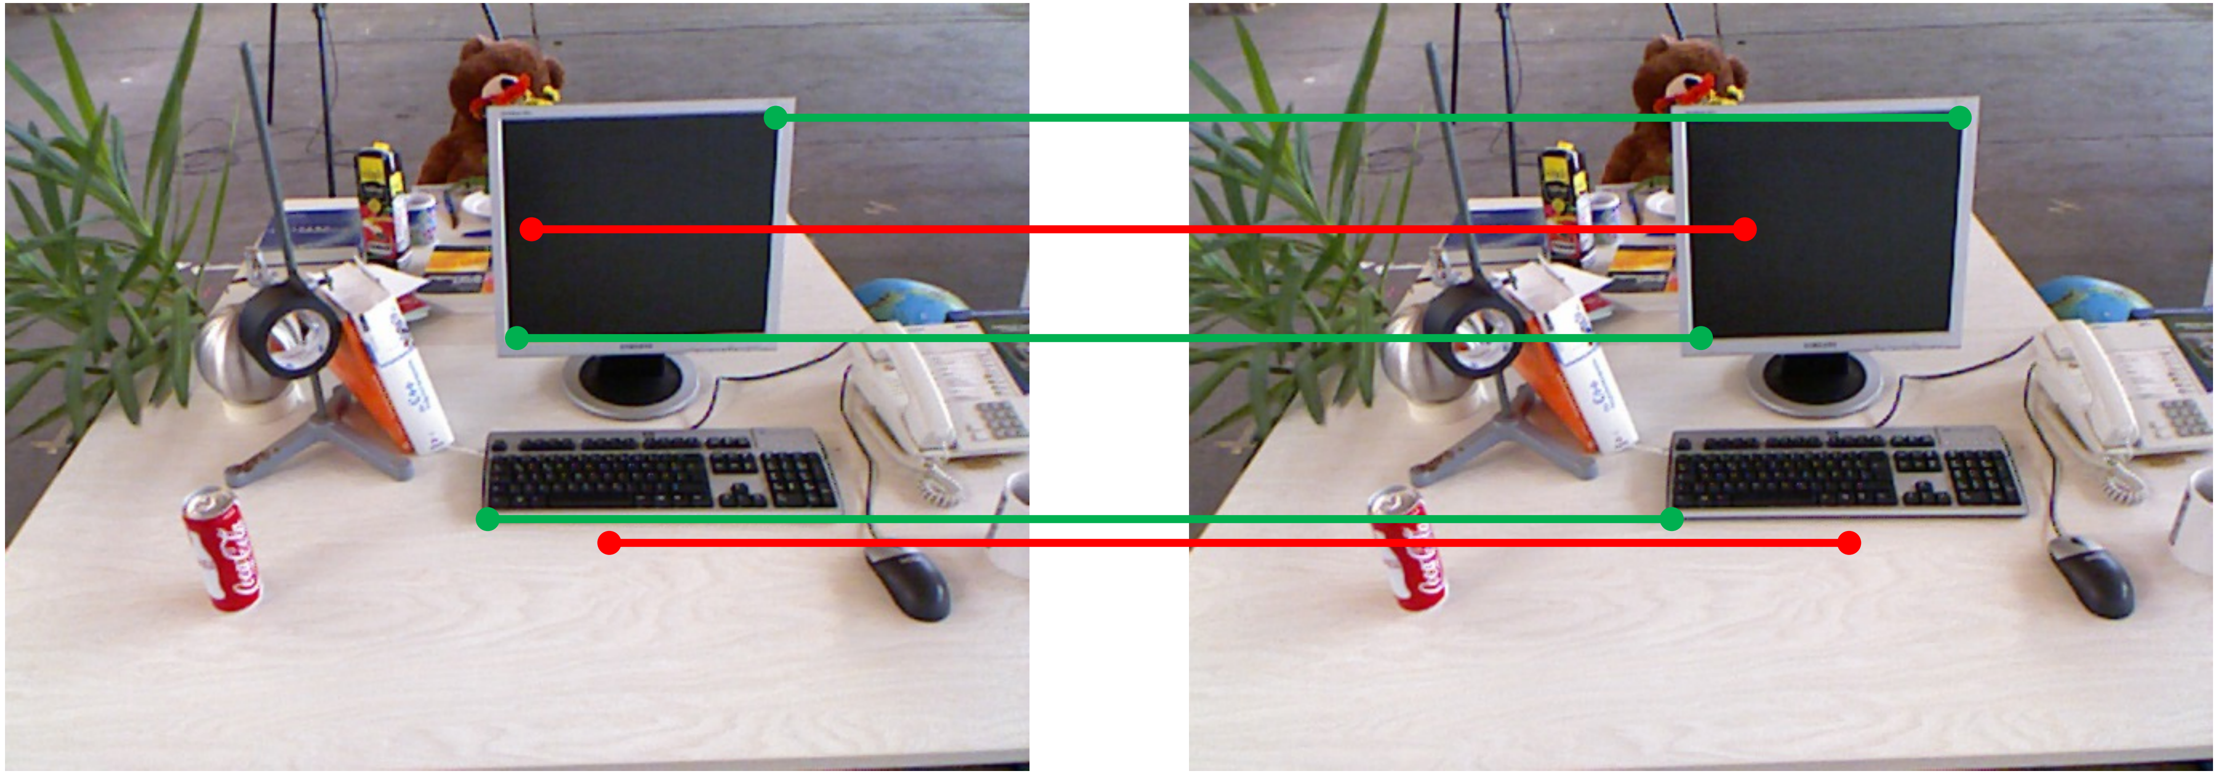

Previous Frame **Current Frame**

Distinguish correct correspondences (**inliers**) from erroneous ones (**outliers**)

Depth Map Estimation for Rigid Scenes

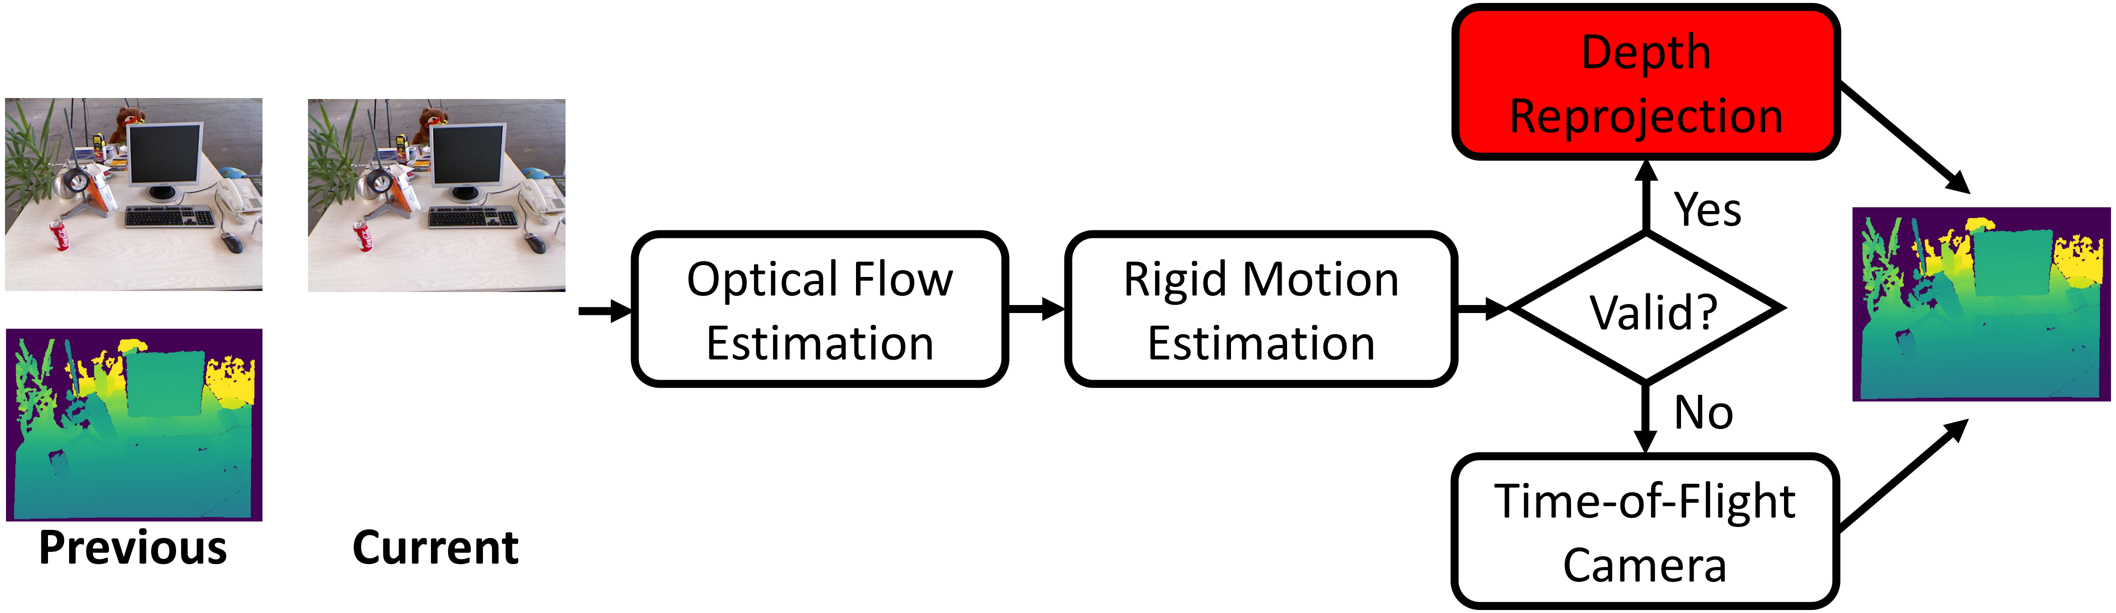

Obtain a new depth map by applying the rigid motion to the previous depth map

Noraky *et al.*, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
Noraky *et al.*, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.

Depth Map Estimation for Rigid Scenes

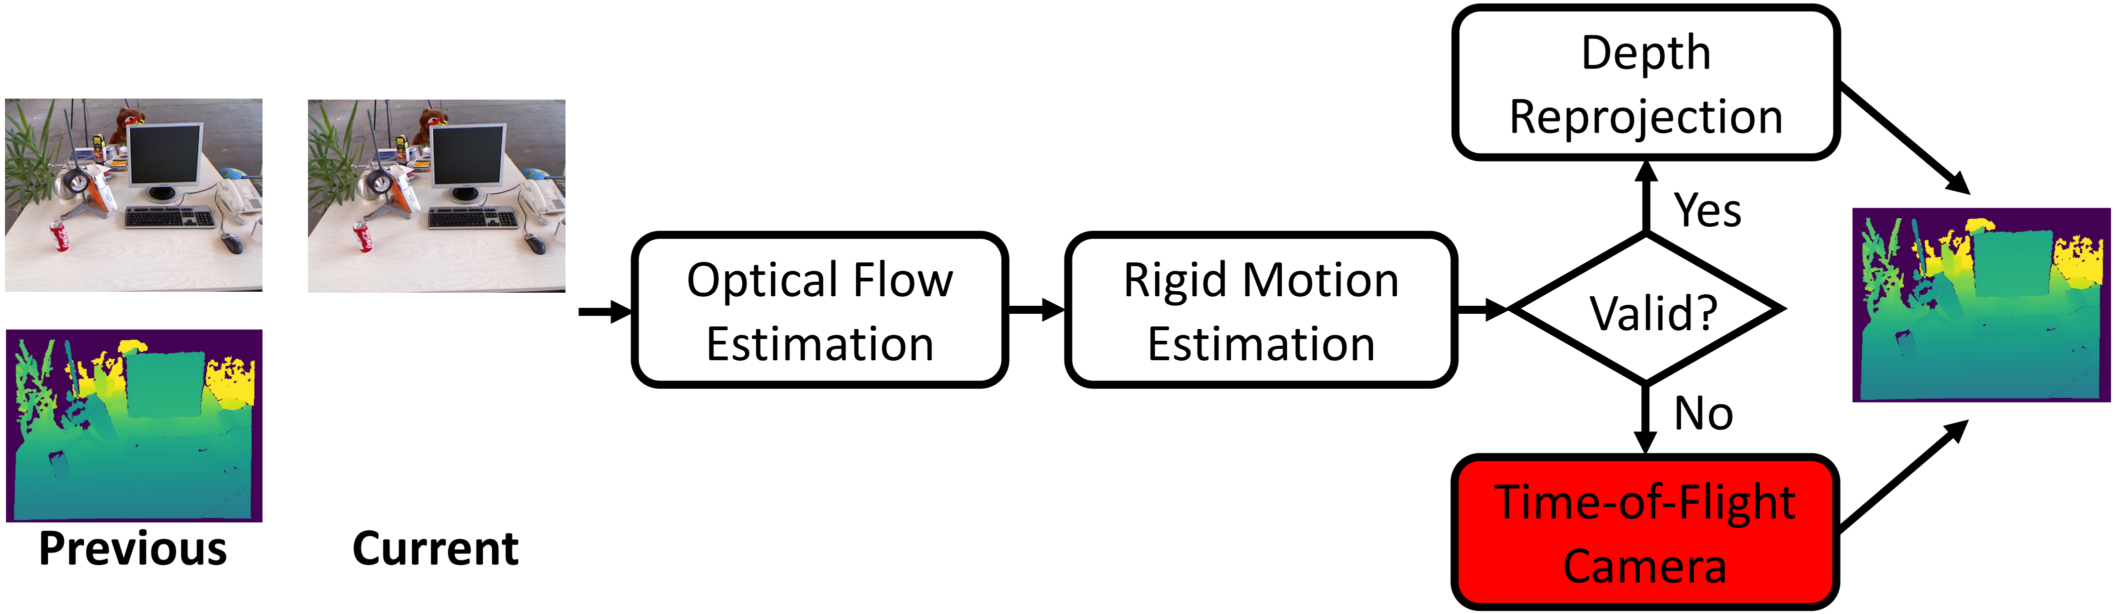

If RANSAC fails, use the time-of-flight camera

Noraky *et al.*, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
Noraky *et al.*, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.

How Does Our Algorithm Perform?

- Evaluate our algorithm using RGB-D datasets: **TUM RGB-D**, **NYU v2**, **Indoor RGB-D**, **CoRBS**, and **ICL-NUIM**
- Estimate depth using the consecutive images and use the depth map in the dataset when the rigid motion cannot be estimated
- Quantify the duty cycle, mean relative error (MRE), and estimation frame rate on the ODROID-XU3 embedded processor

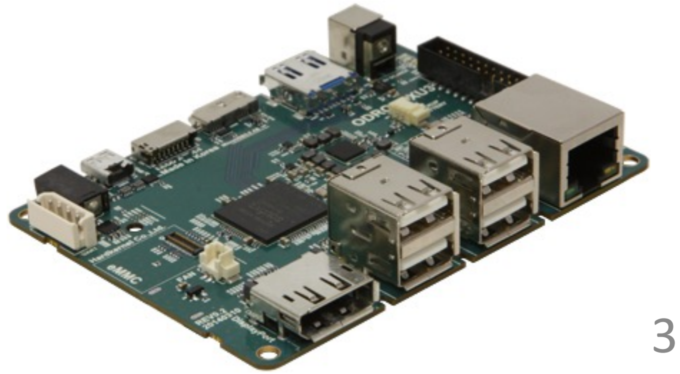

How Low Can We Reduce the Duty Cycle?

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)

Find the lowest duty cycle at which we can maintain a mean relative error of 1%

How Low Can We Reduce the Duty Cycle?

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)
This Work	15.0	0.96	30

Find the lowest duty cycle at which we can maintain a mean relative error of 1%

What Is the Impact of Using Block Matching?

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)
This Work	15.0	0.96	30

Compare to a variant of our approach that computes sub-pixel optical flow (This Work + Sub)

What Is the Impact of Using Block Matching?

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15

What Is the Impact of Using Block Matching?

Approach	**Duty Cycle (%)**	**MRE (%)**	**Frame Rate (FPS)**
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15

Using sub-pixel optical flow **decreases the MRE** but also halves the frame rate

What Is the Impact of Using Block Matching?

Approach	**Duty Cycle (%)**	**MRE (%)**	**Frame Rate (FPS)**
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15

Using sub-pixel optical flow decreases the MRE but also **halves the frame rate**

Balance Accuracy With Estimation Frame Rate

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15

Using sub-pixel optical flow decreases the MRE but also halves the frame rate

What Is the Impact of Using Rigid Motion?

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15

Compare to Wang *et al.*[*], which uses dense optical flow to remap the previous depth map

[*]Wang *et al*., “Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation,” ISCAS, 2010.

Using Rigid Motion Increases Accuracy and Efficiency

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15
Wang *et al.*	15.0	3.20	0.83

*Wang *et al*., "Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation," ISCAS, 2010.

Using Rigid Motion Increases Accuracy and Efficiency

Approach	**Duty Cycle (%)**	**MRE (%)**	**Frame Rate (FPS)**
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15
Wang *et al.*	15.0	3.20	0.83

Reprojecting the previous depth map allows us to account for changes in depth

*Wang *et al*., "Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation," ISCAS, 2010.

Using Rigid Motion Increases Accuracy and Efficiency

Approach	Duty Cycle (%)	MRE (%)	Frame Rate (FPS)
This Work	15.0	0.96	30
This Work + Sub	15.0	0.87	15
Wang *et al.*	15.0	3.20	0.83

Rigid motion can be estimated with sparse optical flow and linear least squares

*Wang *et al*., "Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation," ISCAS, 2010.

What About the System Power?

Category	**Power (W)**
ToF Camera (< 3 m)	1-5

What About the System Power?

Category	Power (W)
ToF Camera (< 3 m)	1-5
ODROID-XU3	0.69
ODROID-XU3 Idle	0.29

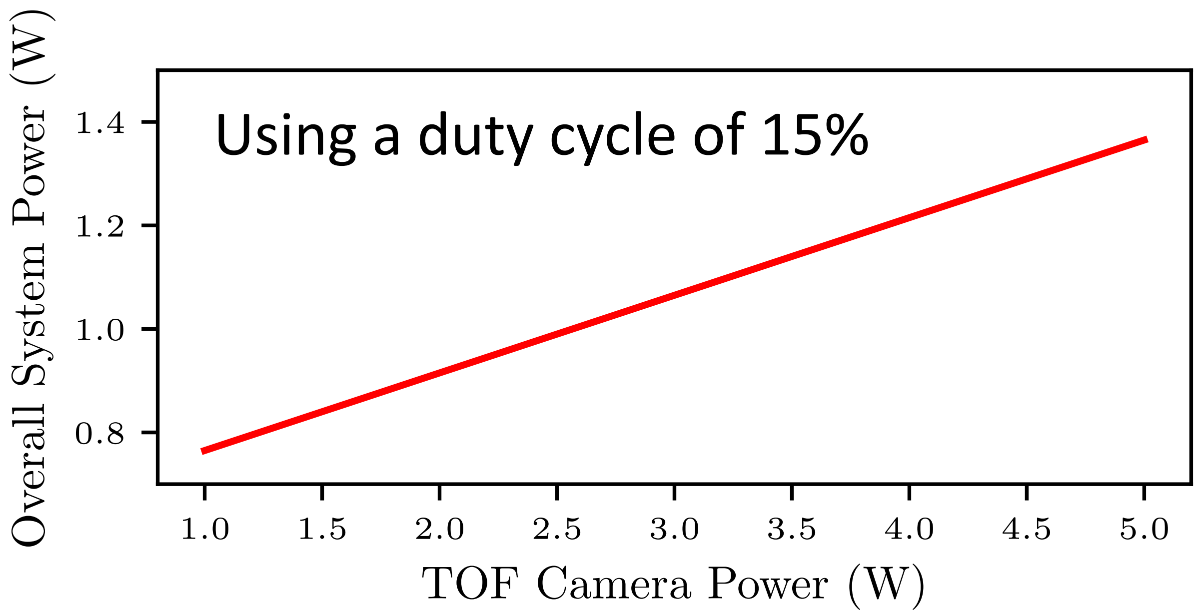

Overall system power reduced by up to 73%

Example of Estimated Depth Maps

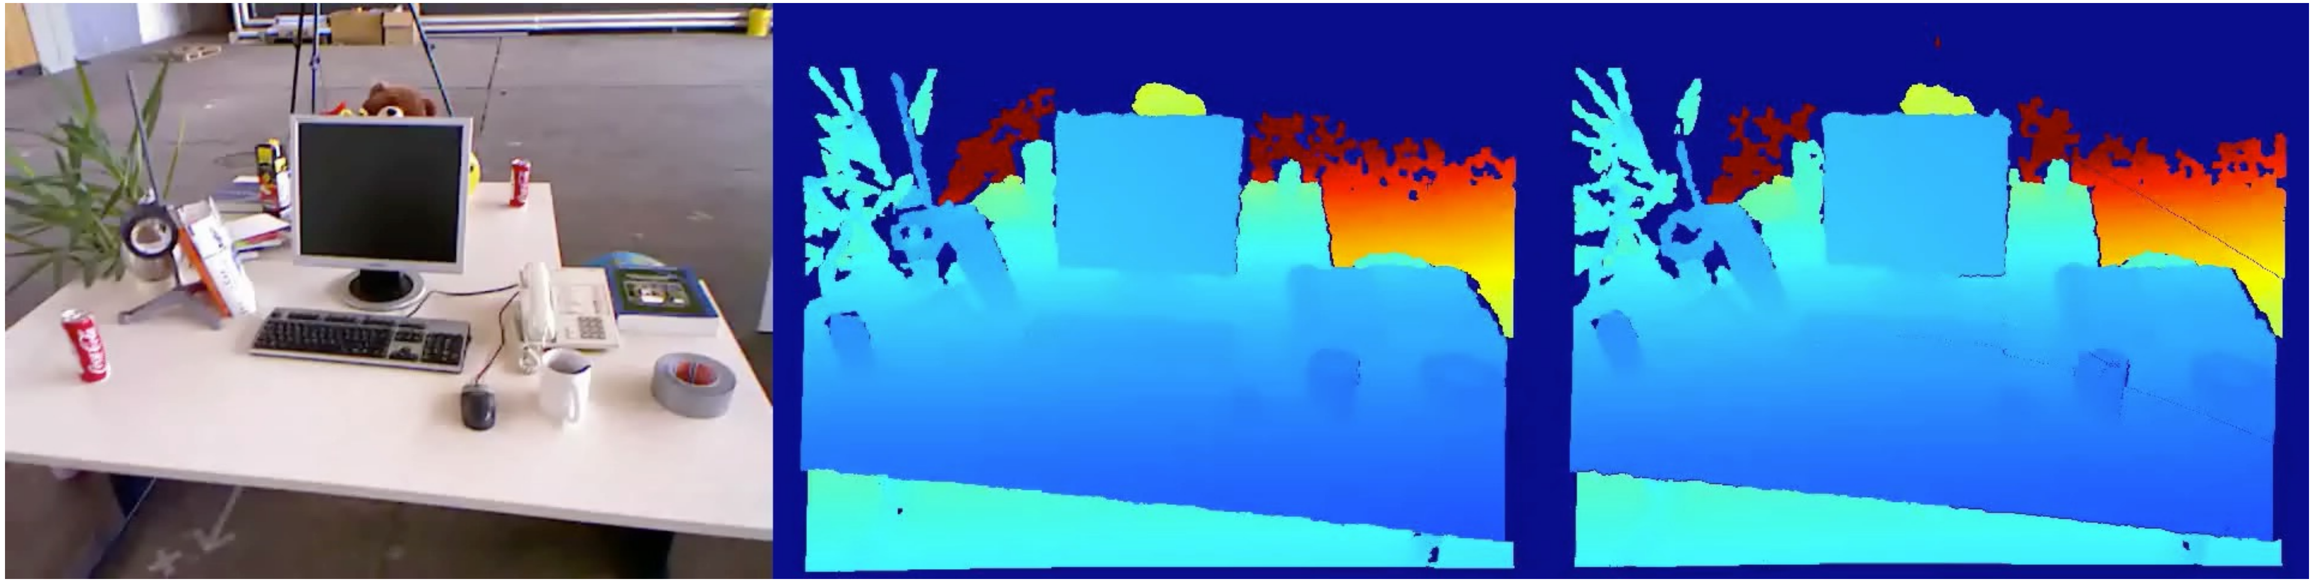

Image **Measured Depth** **Estimated Depth**

Our Depth Maps Can Be Used in Augmented Reality

Image

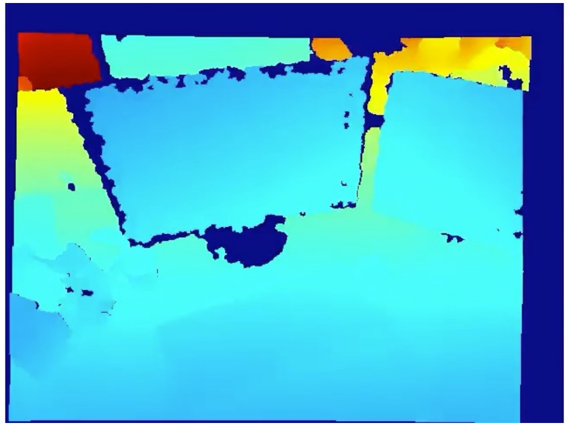

Measured Depth

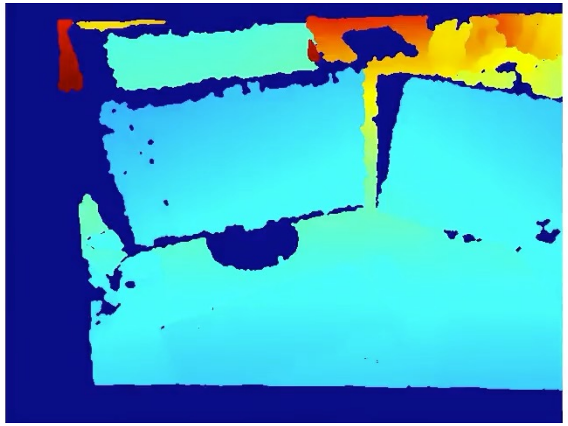

Estimated Depth

Augmented Reality

Joint work with Alan Cheng (SuperUROP + MEng)

Summary of Contributions

- **Key Insight:** We can estimate accurate and dense depth maps efficiently by reprojecting a previous one using the estimated rigid motion
- Using the rigid motion allows us to account for changes in depth, and it can be efficiently estimated using sparse block matching with RANSAC
- We can estimate depth maps in real-time on a low power embedded processor and adaptively control the ToF camera
- Reduce the usage of the ToF camera by up to 85% (and the system power by 73%) while estimating depth within 1% of the ground truth

Our Publications: ICIP 2017, TCSVT 2020

Outline

Strictly Computation | Strictly ToF Camera

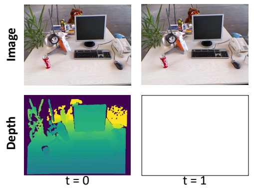

- Reduce the Usage of the ToF Camera
  - Depth Map Estimation for Rigid Scenes
  - **Depth Map Estimation for Dynamic Scenes**
- Reduce the Light the ToF Camera Emits
  - Adaptive Pulse Control
- Summary of Thesis Contributions

Depth Map Estimation for Dynamic Scenes

- Assume that the scene is **locally rigid**
- Many approaches first **segment** the scene into rigid regions and then estimate the depth in each region
- In our work, we cluster the rigid motions and use them to estimate a new depth map

Depth Map Estimation for Dynamic Scenes

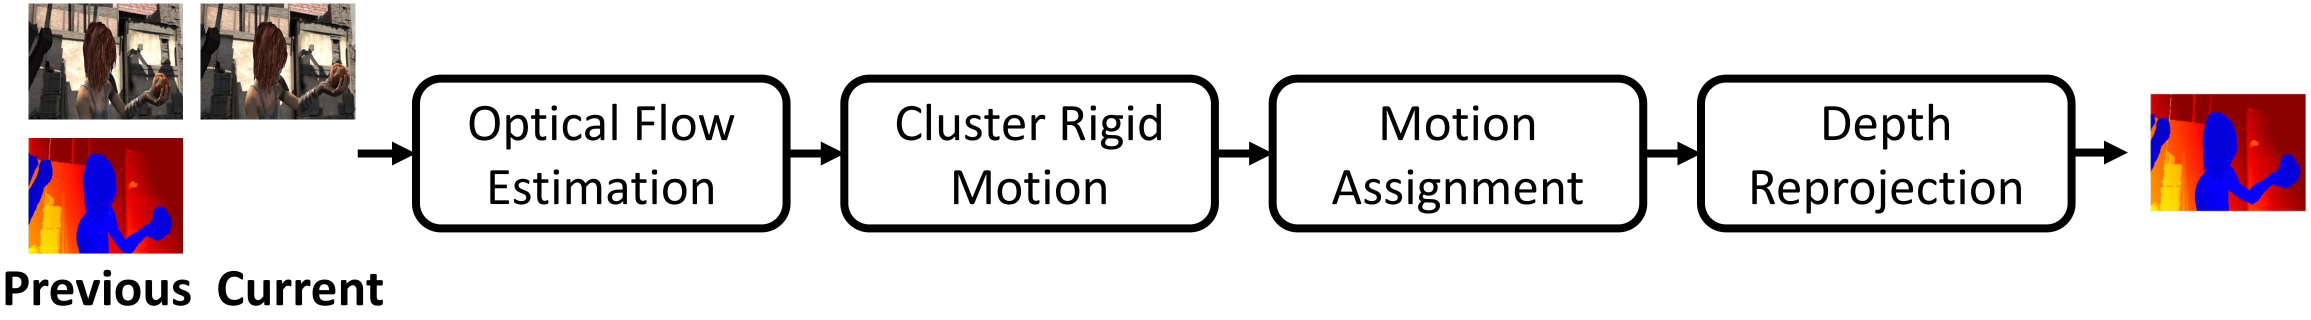

Noraky *et al.*, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.

Depth Map Estimation for Dynamic Scenes

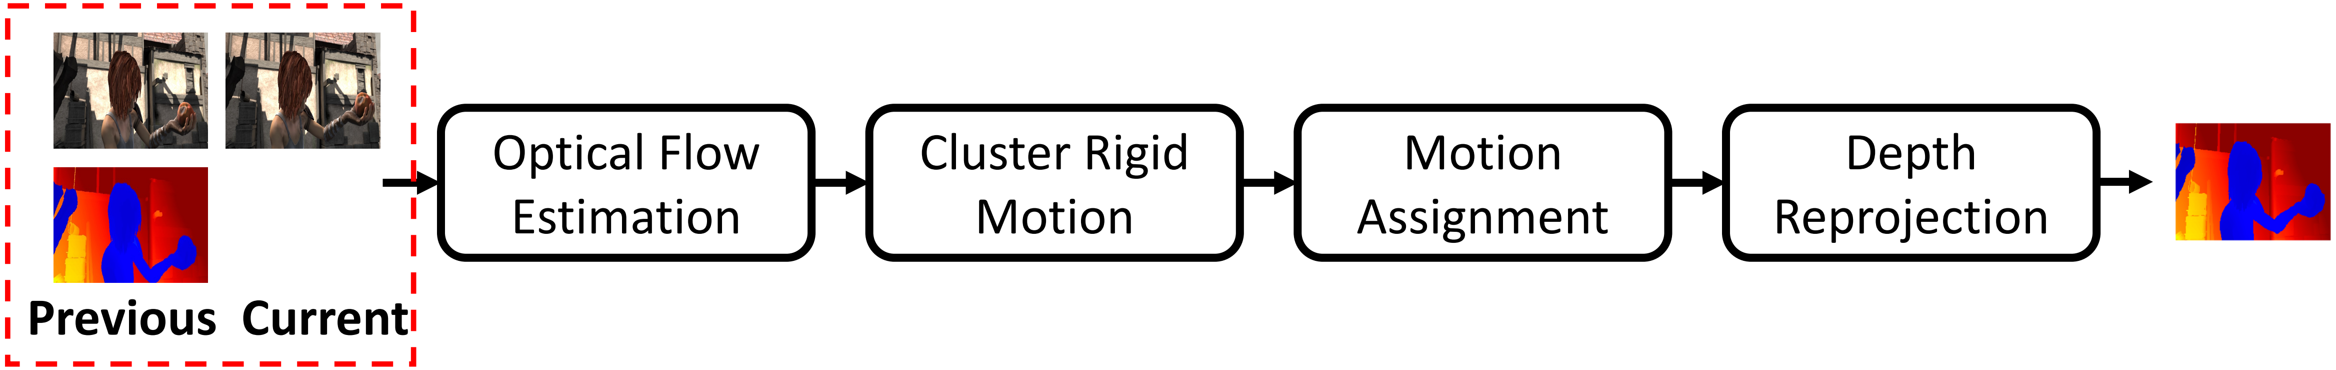

Our inputs are consecutive images and a previously measured depth map

Noraky *et al.*, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.

Depth Map Estimation for Dynamic Scenes

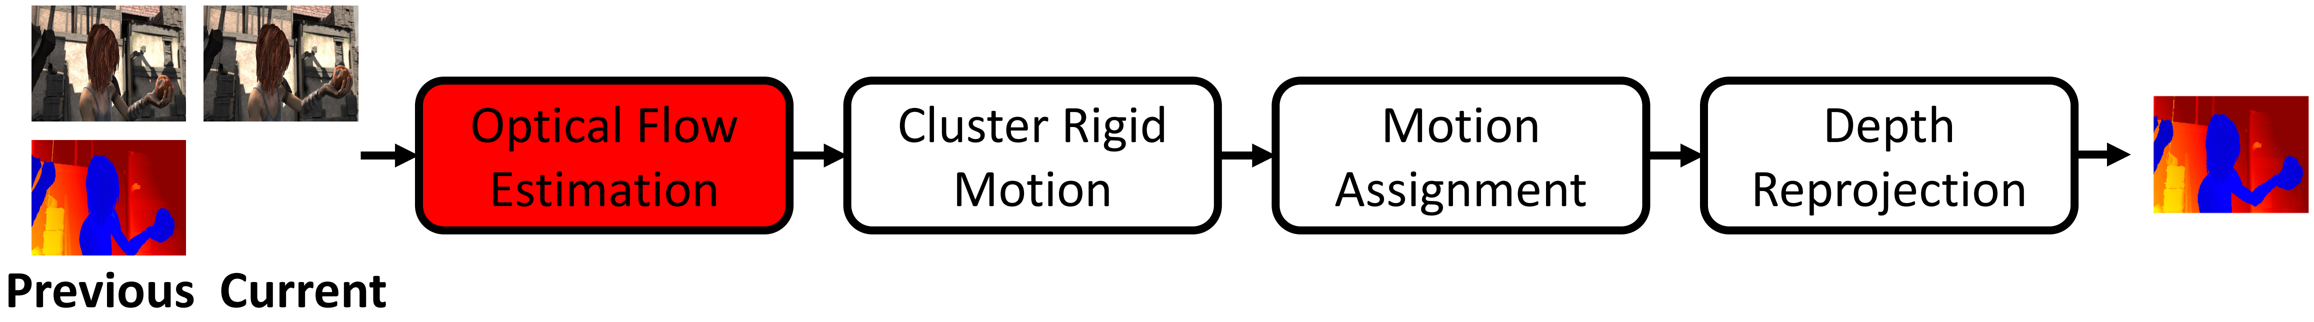

Estimate **sparse** subpixel optical flow at corners

Noraky *et al.*, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.

Depth Map Estimation for Dynamic Scenes

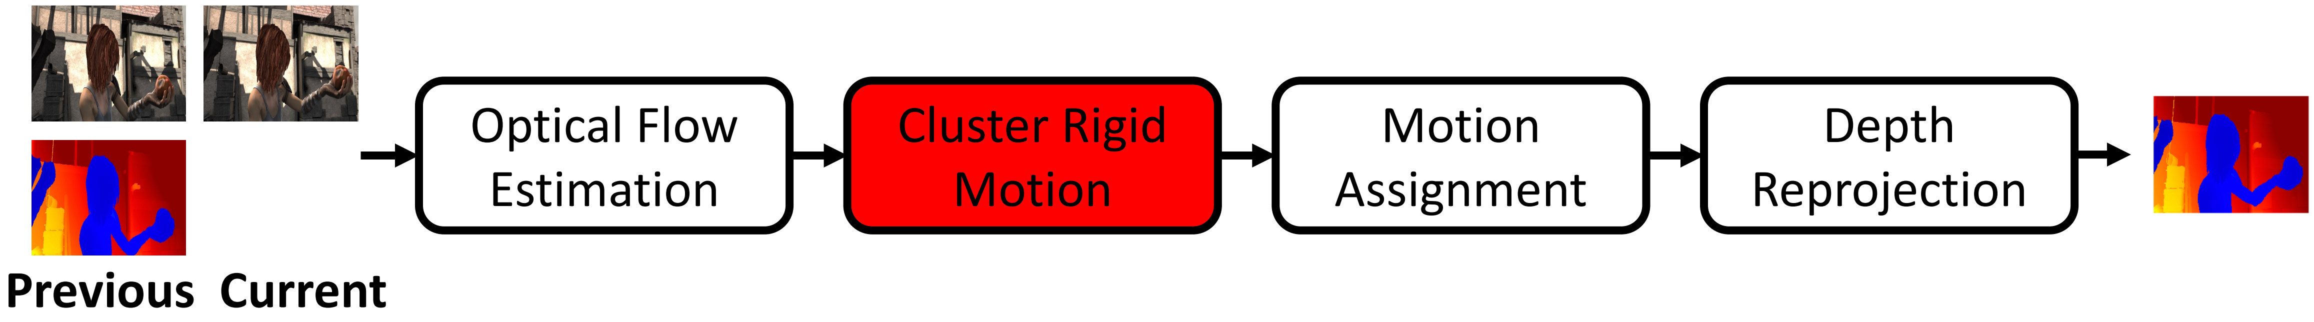

Estimate the rigid motions in the scene by clustering them

Noraky *et al.*, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.

Cluster Rigid Motion

Corners where the optical flow is estimated

Cluster Rigid Motion

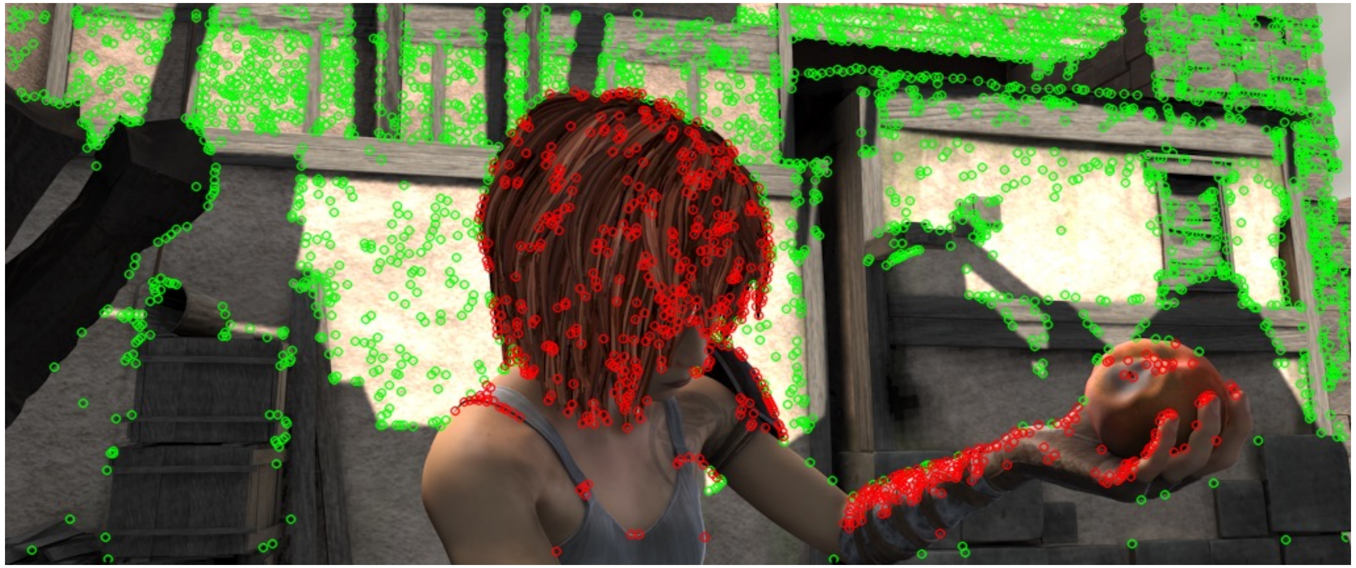

Use RANSAC to estimate the rigid motion and **inliers**

Cluster Rigid Motion

Remove the pixels that correspond to the largest inlier set

Cluster Rigid Motion

Repeat iteratively for the remaining pixels

Depth Map Estimation for Dynamic Scenes

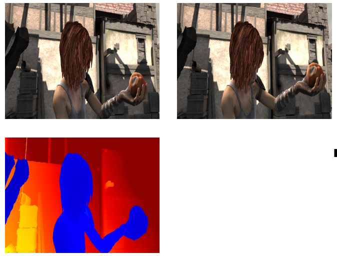

Previous **Current**

Optical Flow Estimation → Cluster Rigid Motion → Motion Assignment → Depth Reprojection →

Assign the pixel motion to a pixel if it minimizes its photometric error

Which rigid motion does a pixel correspond to?

Noraky *et al.*, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.

Computing the Photometric Error

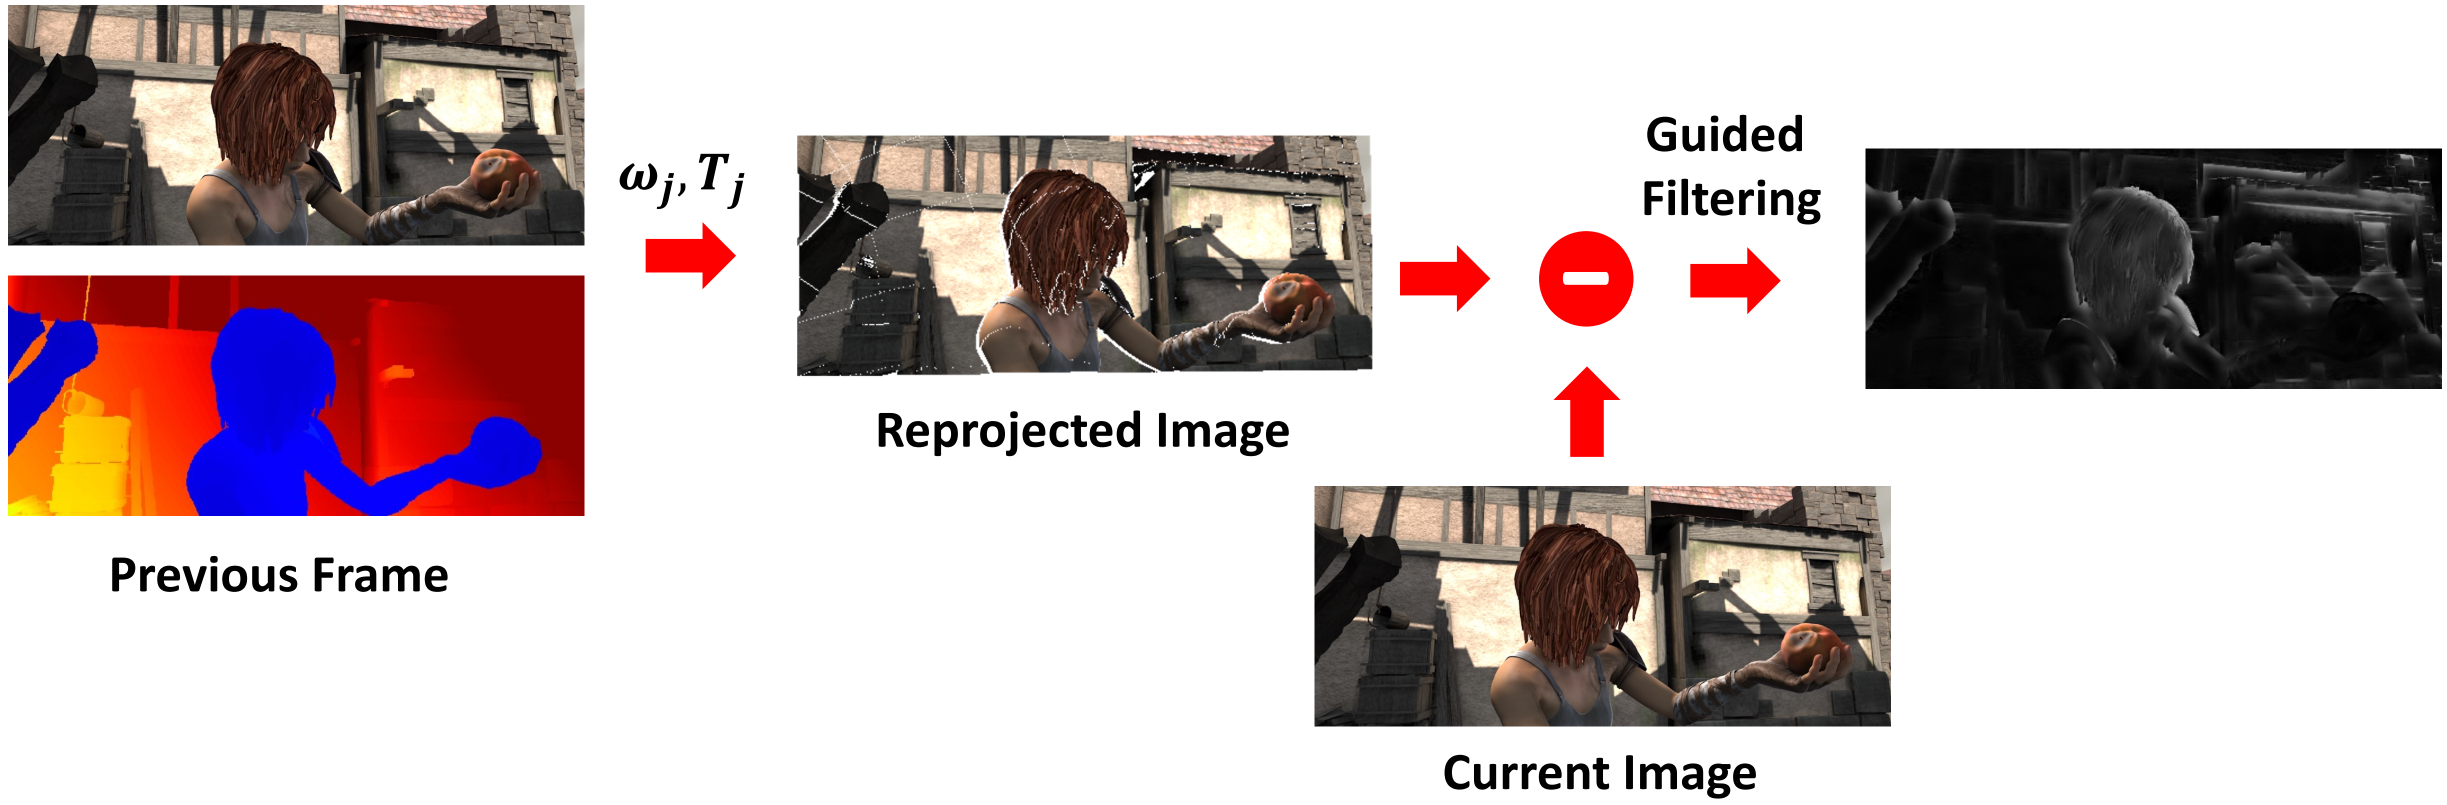

Repeat this process for each of the estimated rigid motions

Use the Photometric Error to Assign the Rigid Motion

The background has low photometric error

Use the Photometric Error to Assign the Rigid Motion

The hand has low photometric error

Use the Photometric Error to Assign the Rigid Motion

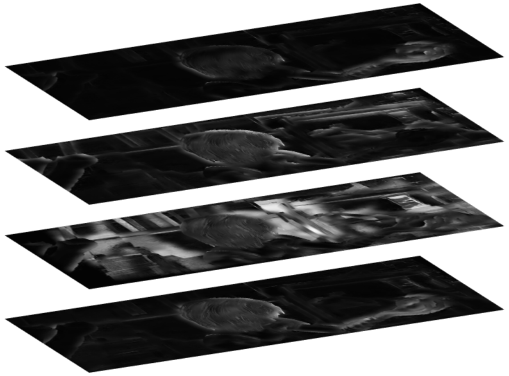

Cost Volume

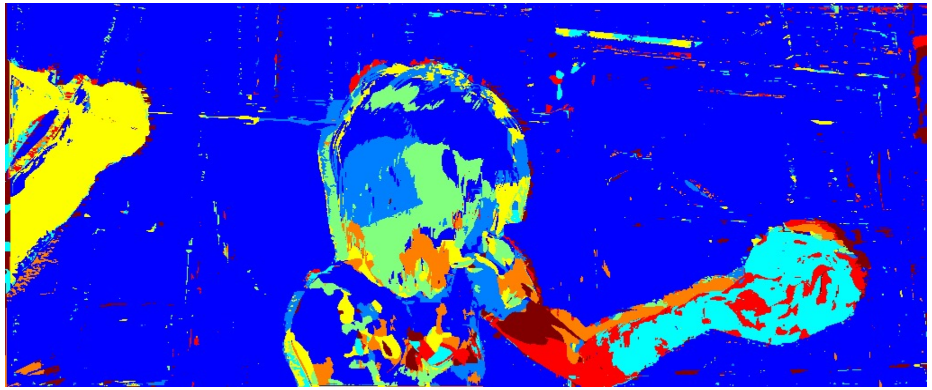

Assigned Rigid Motion

For each pixel, assign the rigid motion that minimizes its photometric error

We can assign the rigid motion without prior segmentation

Depth Map Estimation for Dynamic Scenes

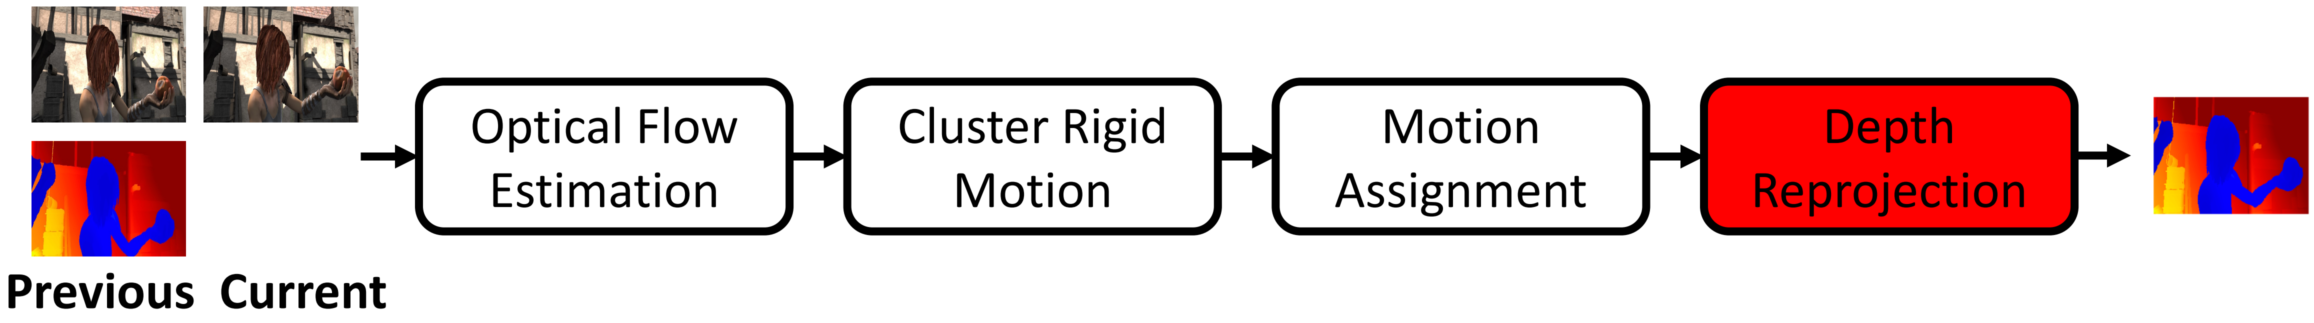

Use the assigned rigid motion to reproject the previous depth map

Noraky *et al.*, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.

How Does Our Algorithm Perform?

- Evaluate our algorithm using RGB-D datasets: **EPFL DS**, **MPI Sintel**, **TUM RGB-D**, and **VKITTI**
- Estimate depth using the consecutive images and use the depth map in the dataset at regular intervals
- Quantify the duty cycle, mean relative error (MRE), and estimation frame rate on a laptop computer

How Low Can We Reduce the Duty Cycle?

Approach	Duty Cycle (%)	MRE (%)

Find the lowest duty cycle at which we can maintain a mean relative error of 1%

How Low Can We Reduce the Duty Cycle?

Approach	Duty Cycle (%)	MRE (%)
This Work	33.3	0.96

Find the lowest duty cycle at which we can maintain a mean relative error of 1%

Dynamic Scenes Are Challenging

Approach	Duty Cycle (%)	MRE (%)
This Work	33.3	0.96

We use the ToF camera more than **twice as much** as before (duty cycle of 33.3% vs 15.0%)

Dynamic Scenes Are Challenging

Approach	Duty Cycle (%)	MRE (%)
This Work	33.3	0.96
This Work	15.0	1.74

MRE **increases by 81%** when estimating at the same duty cycle (MRE of 1.74% vs 0.96%)

What Is the Impact of Using Rigid Motions?

Approach	**Duty Cycle (%)**	**MRE (%)**
This Work	33.3	0.96
This Work	15.0	1.74

Compare to Wang *et al.*[*], which uses dense optical flow to remap previous depth map

[*]Wang *et al*., "Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation," ISCAS, 2010.

What Is the Impact of Using Rigid Motions?

Approach	Duty Cycle (%)	MRE (%)
This Work	33.3	0.96
This Work	15.0	1.74
Wang *et al.*	33.3	2.01
Wang *et al.*	15.0	5.14

Compare to Wang *et al.**, which uses dense optical flow to remap previous depth map

*Wang *et al*., "Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation," ISCAS, 2010.

Accounting for Changes in Depth Increases Accuracy

Approach	**Duty Cycle (%)**	**MRE (%)**
This Work	33.3	0.96
This Work	15.0	1.74
Wang *et al.*	33.3	2.01
Wang *et al.*	15.0	5.14

Compare to Wang *et al.*[*], which uses dense optical flow to remap previous depth map

[*]Wang *et al.*, "Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation," ISCAS, 2010.

Accounting for Changes in Depth Increases Accuracy

Approach	**Duty Cycle (%)**	**MRE (%)**
This Work	33.3	0.96
This Work	15.0	1.74
Wang *et al.*	33.3	2.01
Wang *et al.*	15.0	5.14

Compare to Wang *et al.*[*], which uses dense optical flow to remap previous depth map

[*]Wang *et al*., “Depth Maps Interpolation from Existing Pairs of Keyframes and Depth Maps for 3D Video Generation,” ISCAS, 2010.

Do We Need Previous Depth Maps to Begin With?

Approach	**Duty Cycle (%)**	**MRE (%)**
This Work	33.3	0.96
This Work	15.0	1.74
Wang *et al.*	33.3	2.01
Wang *et al.*	15.0	5.14

Compare to Kumar *et al.*[*], which estimates depth maps using only RGB images

[*]Kumar *et al*., "Monocular Dense 3D Reconstruction of a Complex Dynamic Scene from Two Perspective Frames," ICCV, 2017.

Using Previous Depth Map Increases Accuracy

Approach	**Duty Cycle (%)**	**MRE (%)**
This Work	33.3	0.96
This Work	15.0	1.74
Wang *et al.*	33.3	2.01
Wang *et al.*	15.0	5.14
Kumar *et al.*	0	10.65

Balancing the sensor usage with computation increases accuracy

*Kumar *et al.*, “Monocular Dense 3D Reconstruction of a Complex Dynamic Scene from Two Perspective Frames,” ICCV, 2017.

We Estimate Dense Depth Maps in Near Real-Time

Resolution	This Work (FPS)	Wang *et al.* (FPS)	Kumar *et al.* (FPS)
640x480	33.0	7.0	--
1024x436	12.0	4.3	< 0.0002

By clustering the rigid motion and assigning them using the photometric error, we increase the estimation frame rate and obtain accurate depth maps

Summary of Contributions

- **Key Insight**: We can obtain an accurate depth map efficiently by estimating and assigning the rigid motions in the scene without prior segmentation
- The rigid motions can be clustered using sparse optical flow, increasing the frame rate at which we estimate depth
- By reprojecting the previous image to obtain the photometric error, we can accurately and efficiently assign the rigid motion
- Reduce the usage of the ToF camera by 85% while still estimating dense depth maps within 1.74% of the ground truth in up to real-time

Our Publication: Under Review, 2020

Future Directions

- **Lack of Texture in the Scene**: Need texture to estimate accurate optical flow
  - Explore how the photometric error can be used as a confidence map and to control the ToF camera

- **Missing Depth**: Regions that are uncovered do not have depth
  - Explore low cost infilling methods that can be used

- **Run Tin** ap esc im te decreas
  - Expl ith ate d h it the para roa

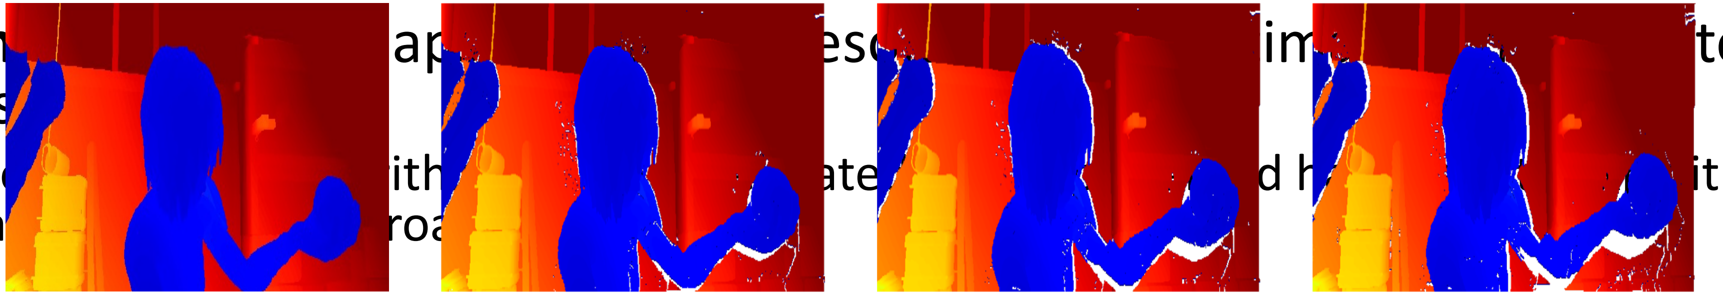

Outline

Strictly Computation | Strictly ToF Camera

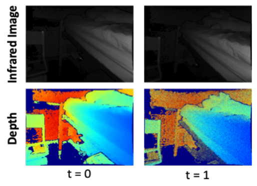

- Reduce the Usage of the ToF Camera
  - Depth Map Estimation for Rigid Scenes
  - Depth Map Estimation for Dynamic Scenes
- Reduce the Light the ToF Camera Emits
  - **Adaptive Pulse Control**
- Summary of Thesis Contributions

What If We Only Want to Use the ToF Camera?

- **Simple Solution:** Reduce the pulses of light the ToF camera emits
- **Reduced Range**: Reflected light cannot be discerned from the ambient light
- **Reduced Depth Resolution**: Depth variance is inversely proportional to the reflected intensity

Noisy Depth Maps Obscure Features

Infrared Image

Regular Power

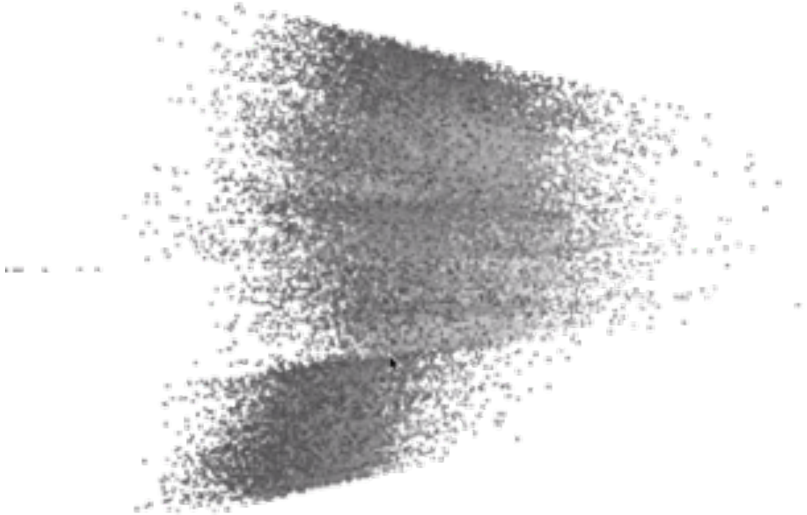

Low Power
(10x Less Light)

Real data captured using ADI ToF Camera*

*AD-96TOF1-EBZ

Noisy Depth Maps Obscure Features

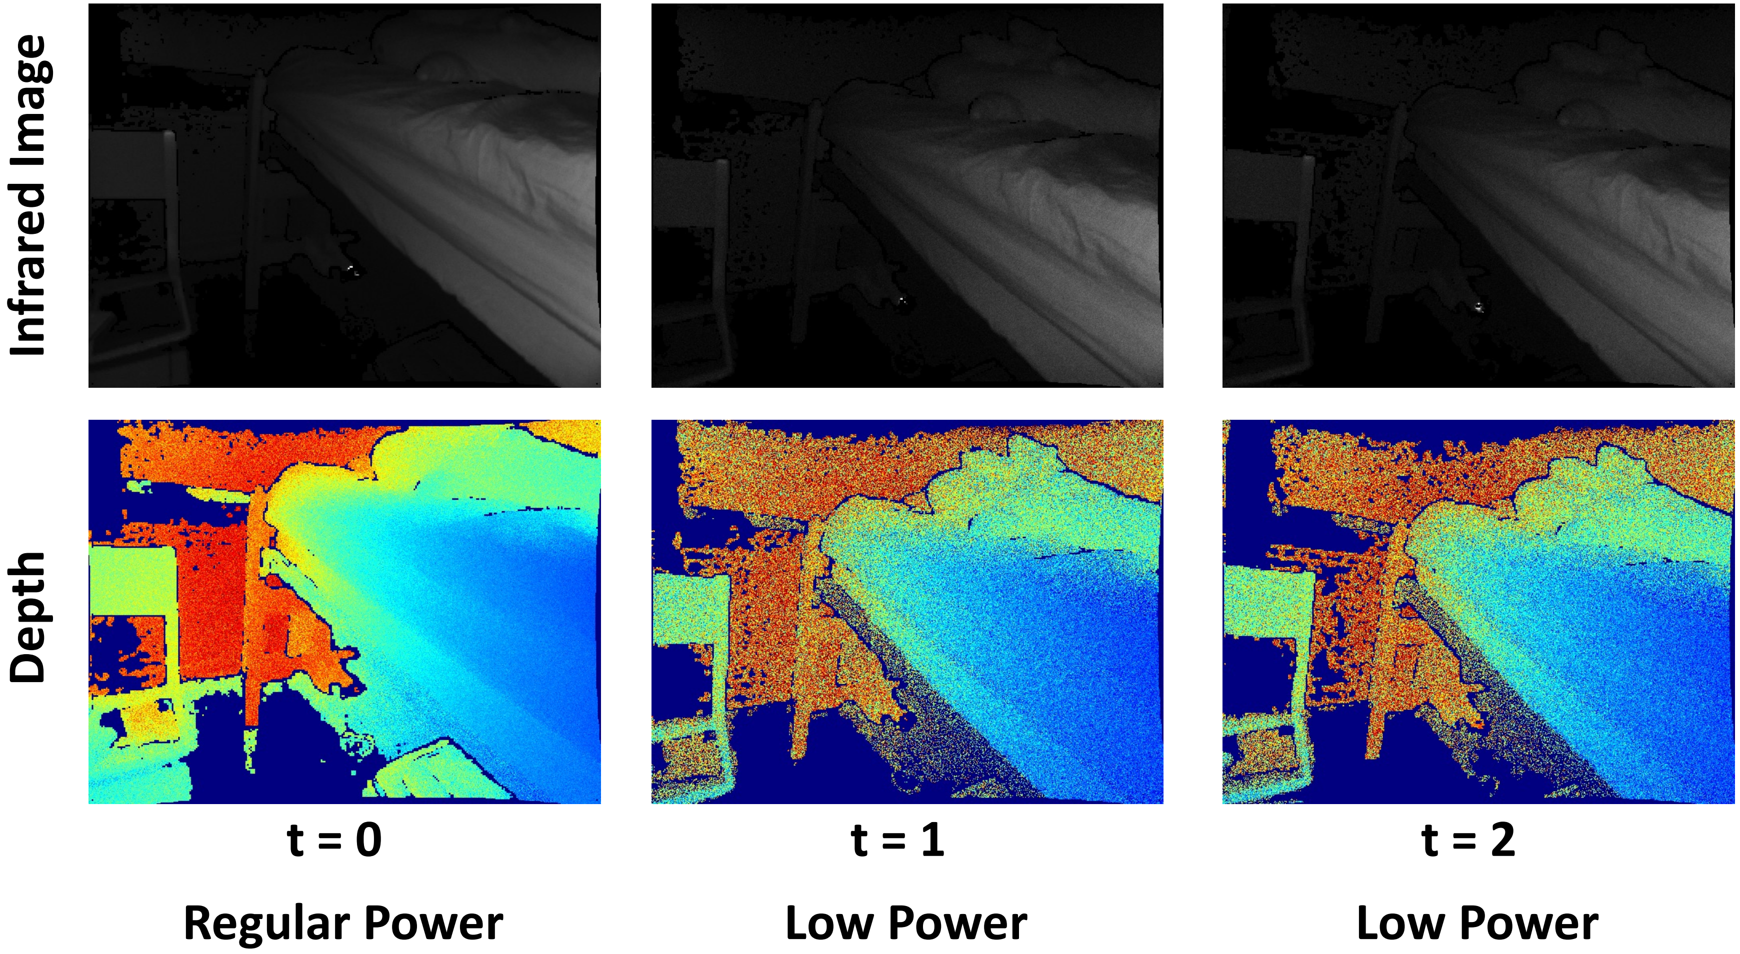

Features in Infrared Images Are Preserved

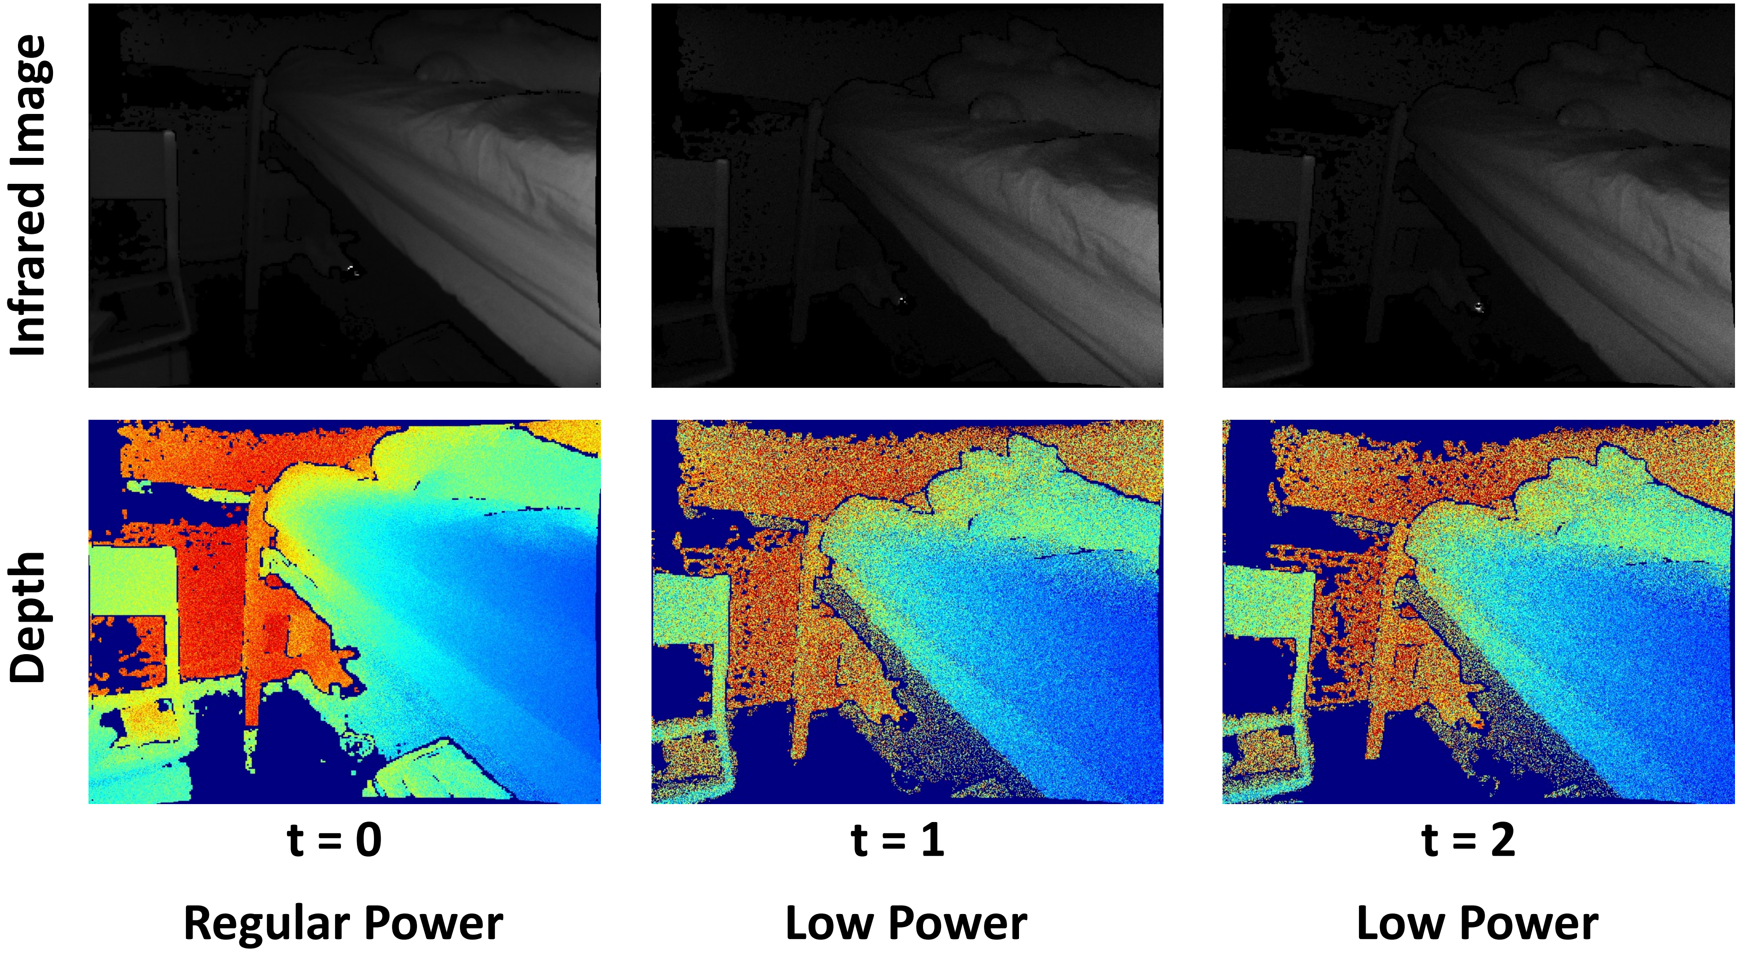

Features in Infrared Images Are Preserved

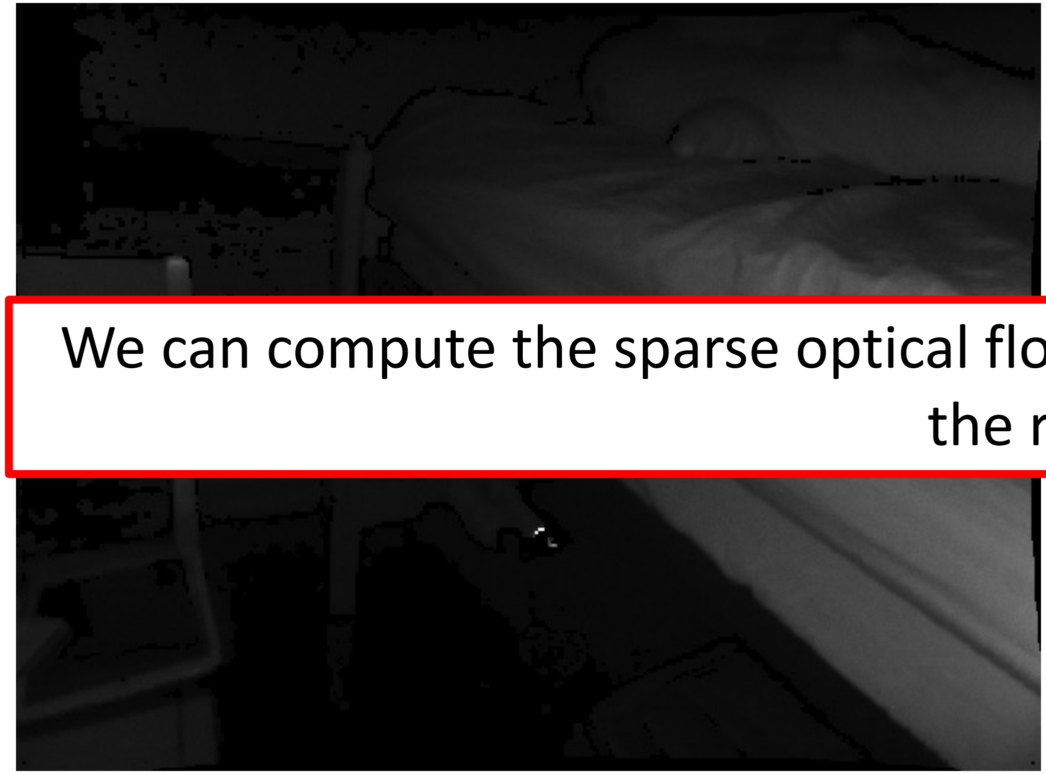

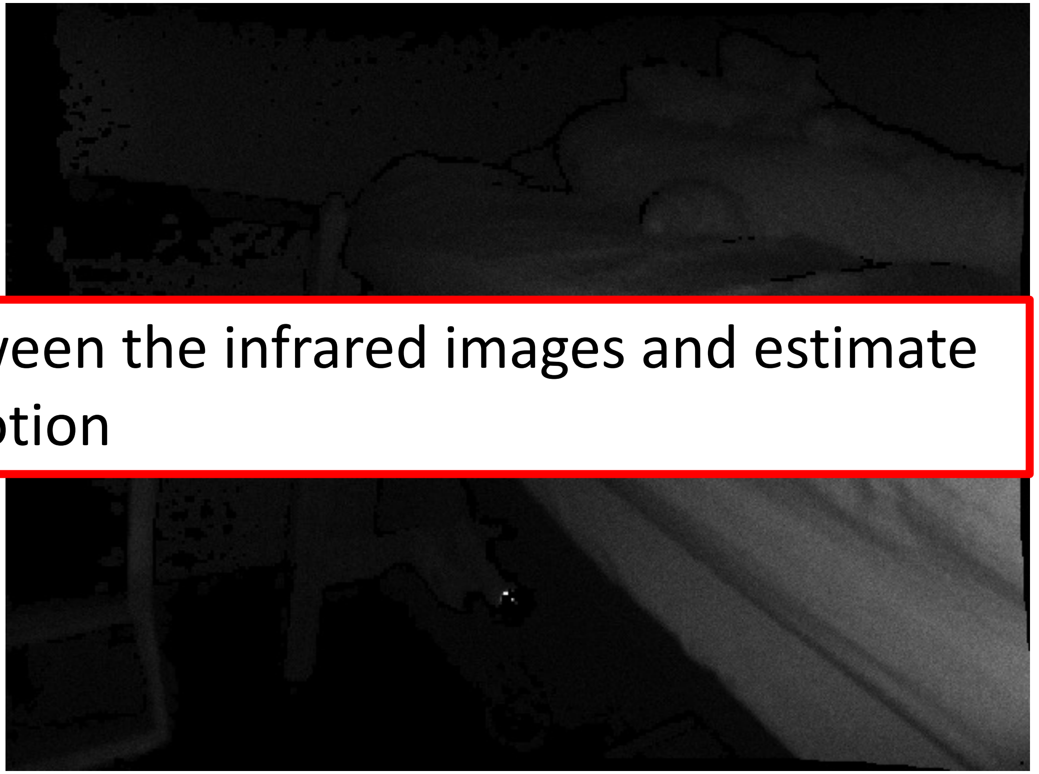

We can compute the sparse optical flow between the infrared images and estimate the rigid motion

t = 0

Regular Power

t = 1

Low Power

Estimate Rigid Motion With Infrared Images

- Use the rigid motion to combine depth maps across frames
- Infrequently obtain regular power depth maps and use them to denoise subsequent low power depth maps
- Difference from Previous Approach:
  - ToF camera is always on, but less light is emitted – Goal is to mitigate noise
  - Use only data from the ToF camera

Adaptive Pulse Control

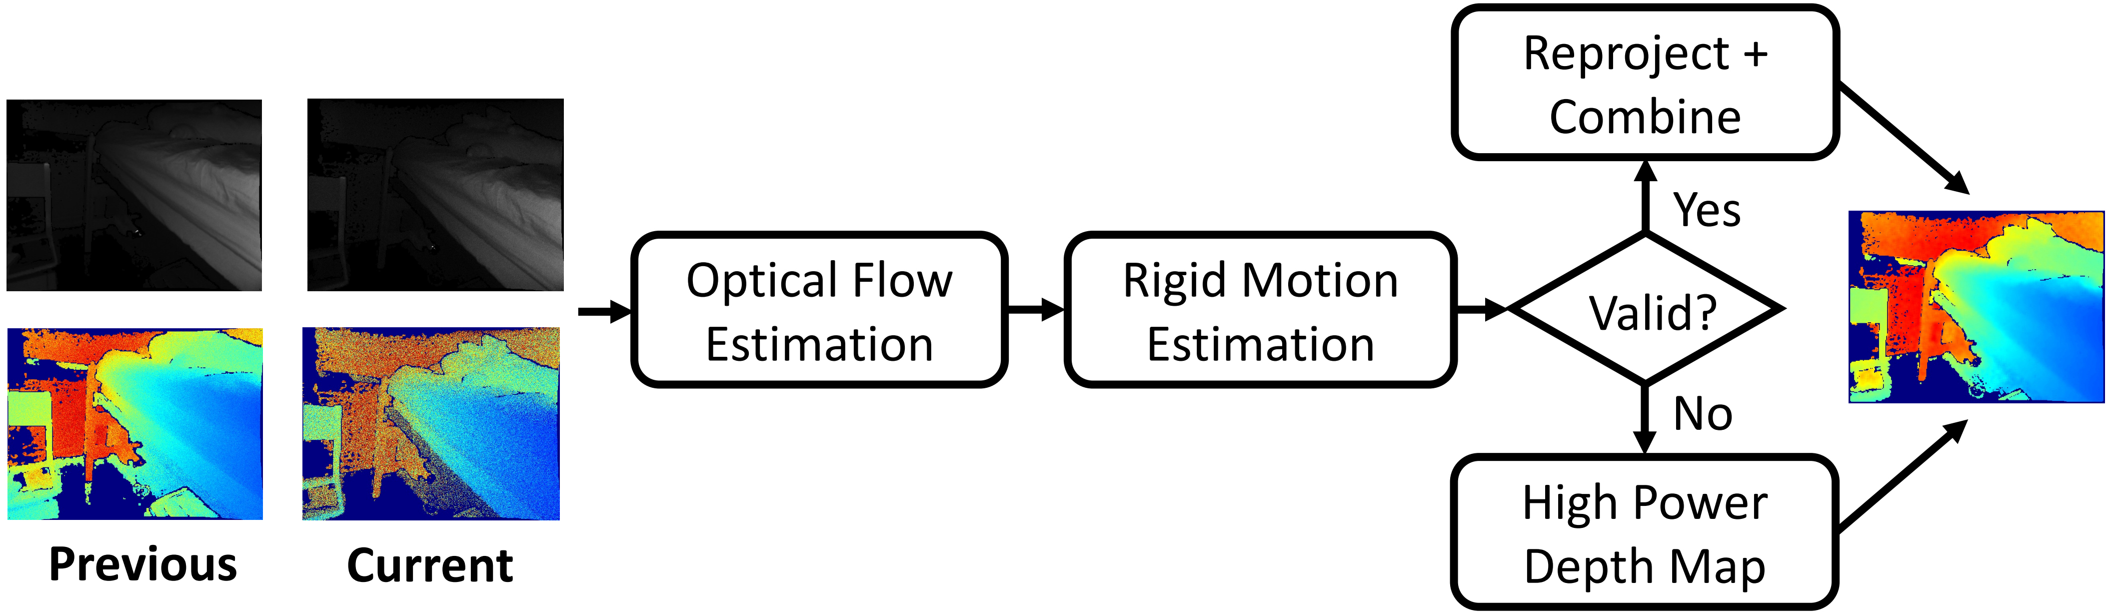

Noraky *et al*., "Low Power Adaptive Time-of-Flight Imaging For Multiple Rigid Objects," ICIP, 2019.

Adaptive Pulse Control

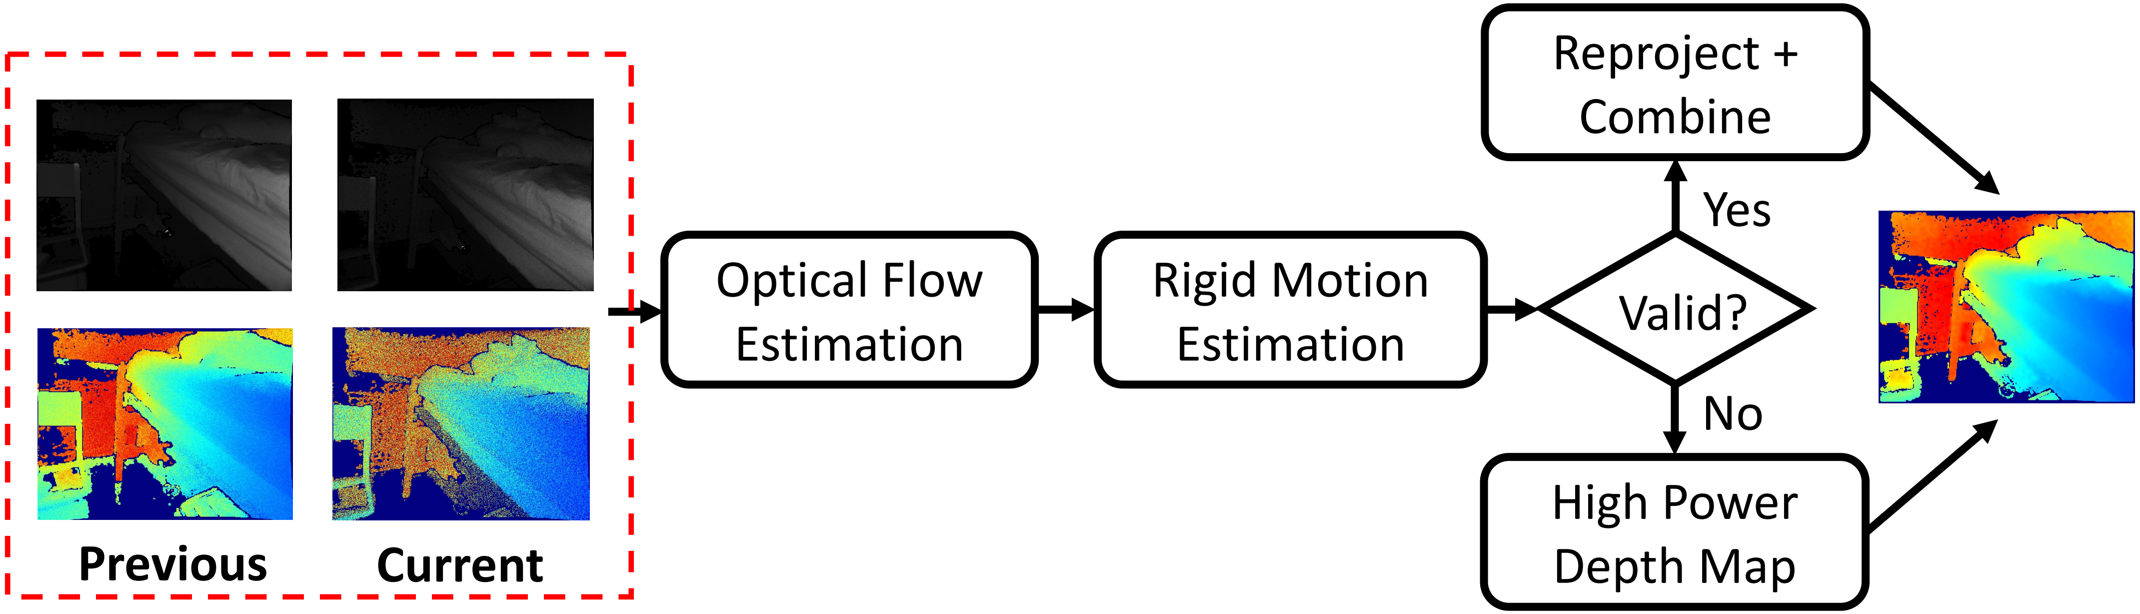

Inputs are the previously filtered frame and the current low power one

Noraky *et al.*, ”Low Power Adaptive Time-of-Flight Imaging For Multiple Rigid Objects," ICIP, 2019.

Adaptive Pulse Control

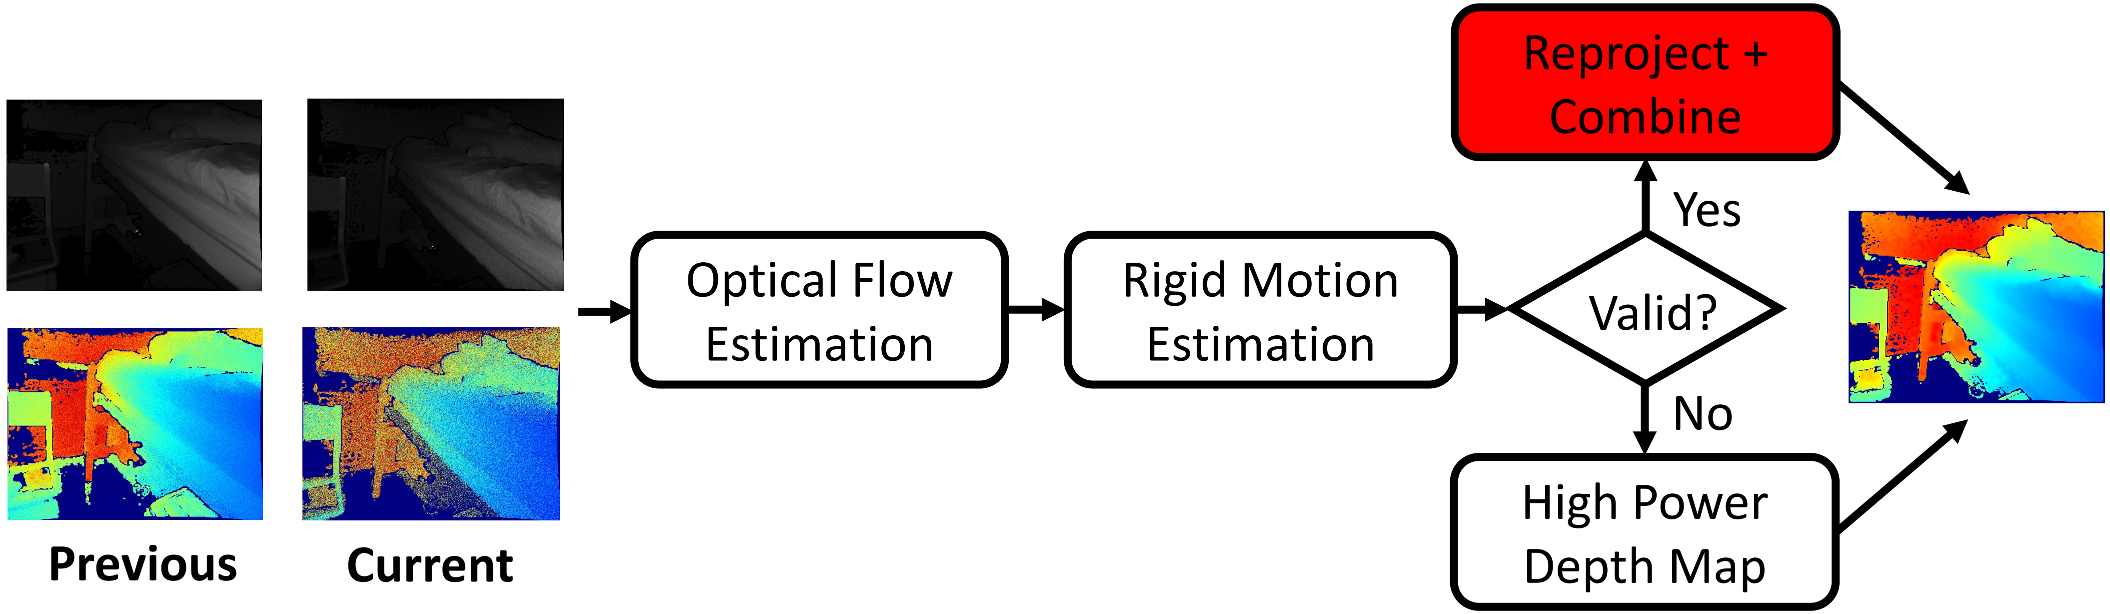

Combine the reprojected depth map with the low power one using a weighted average

Noraky *et al.*, ”Low Power Adaptive Time-of-Flight Imaging For Multiple Rigid Objects," ICIP, 2019.

Adaptive Pulse Control

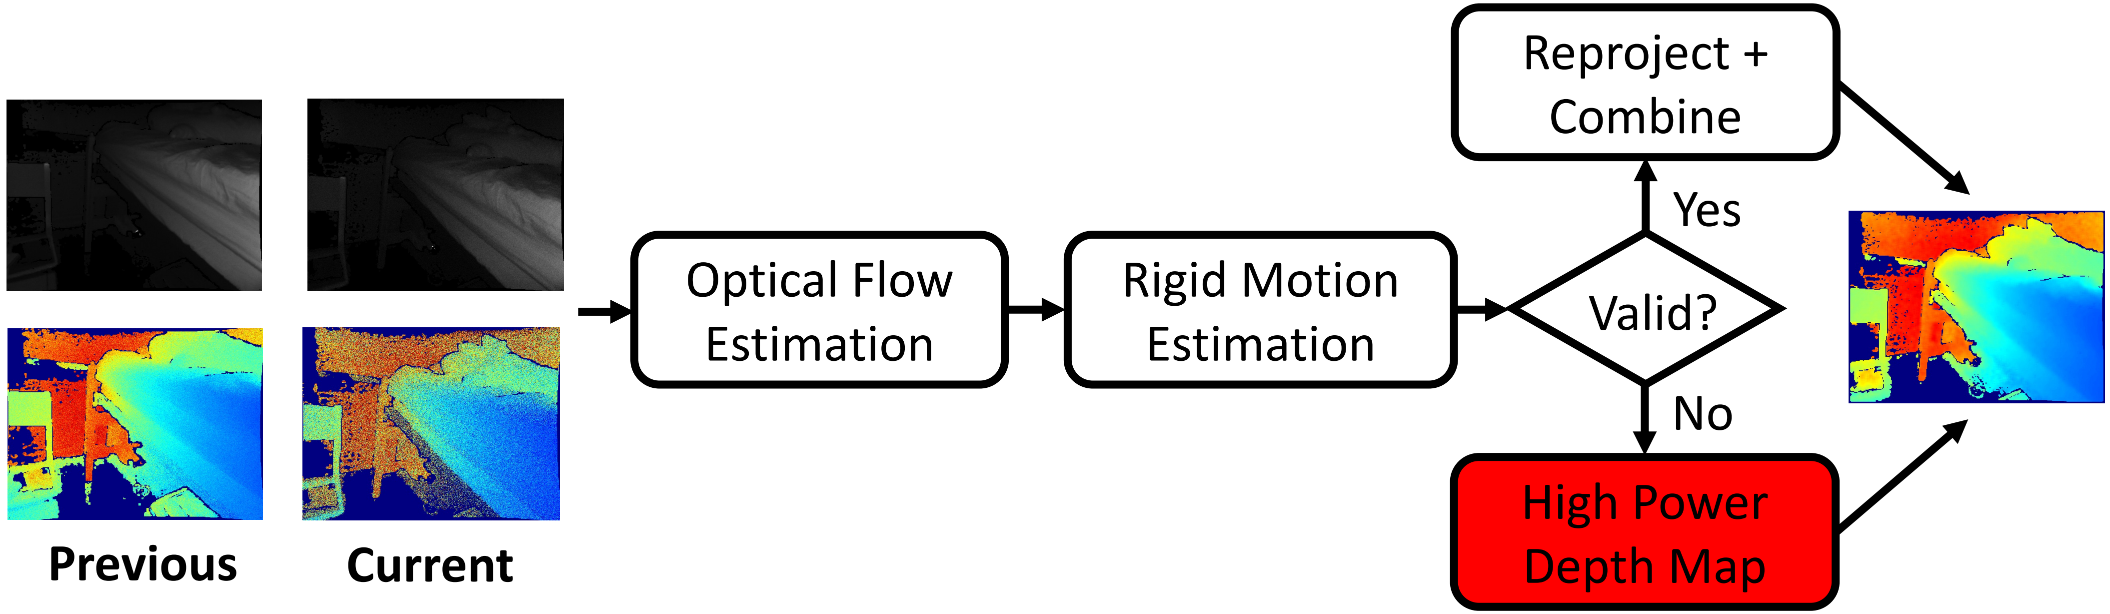

Obtain a regular “high power” power depth map

Noraky *et al*., ”Low Power Adaptive Time-of-Flight Imaging For Multiple Rigid Objects," ICIP, 2019.

How Does Our Algorithm Perform?

- Collect dataset using the Pico Zense DCAM100 ToF camera of common scenes
- Add shot n p e
- Quantify t p a r e
  estimation n U

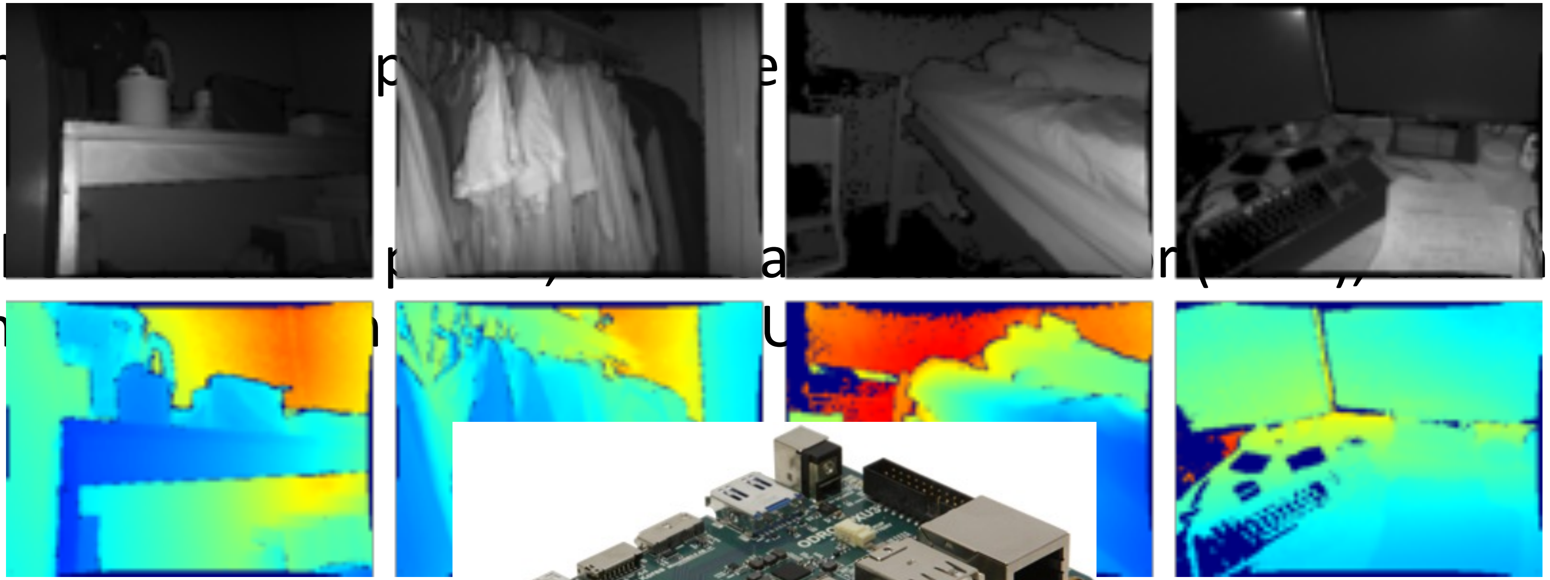

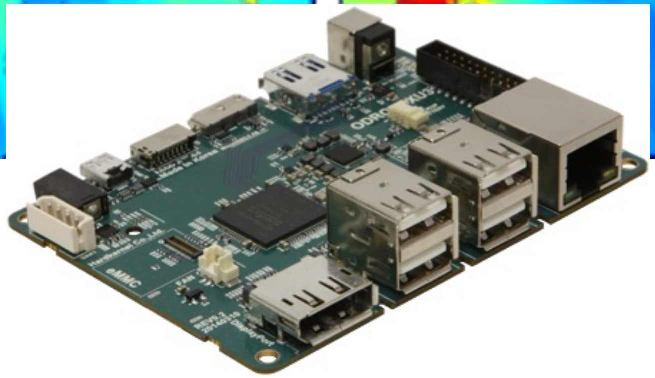

Baseline Results

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30

ToF cameras consume a lot of power but obtain accurate depth with low latency

Baseline Results

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30

Lowering the power increases the MRE significantly

Apply Bilateral Filter to the Depth Maps

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30

Apply Bilateral Filter to the Depth Maps

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6

At Low Powers, Bilateral Filters Are Ineffective

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6

At Low Powers, Bilateral Filters Are Ineffective

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6

The estimation frame rate also decreases

How Does Our Approach Perform?

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6

Using Regular Power Depth Maps Increases Accuracy

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6
This Work	0.19	3.2%	30

Using Regular Power Depth Maps Increases Accuracy

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6
This Work	0.19	3.2%	30

Lower the mean relative error of the low power depth maps by 64% in real-time

Can We Get the Same Accuracy by Increasing Power?

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6
This Work	0.19	3.2%	30

Increase the amount of light the ToF camera emits per frame

Equivalent Power Depth Map Has Higher MRE

Approach	Normalized Power	MRE (%)	Frame Rate (FPS)
Regular Power	1	2.6%	30
Low Power	0.1	8.8%	30
Bilateral Filter	0.1	6.3%	8.6
This Work	0.19	3.2%	30
Equivalent Power	0.19	6.2%	30

At low power, shot noise is especially pronounced

Visualizing the Impact of Our Algorithm

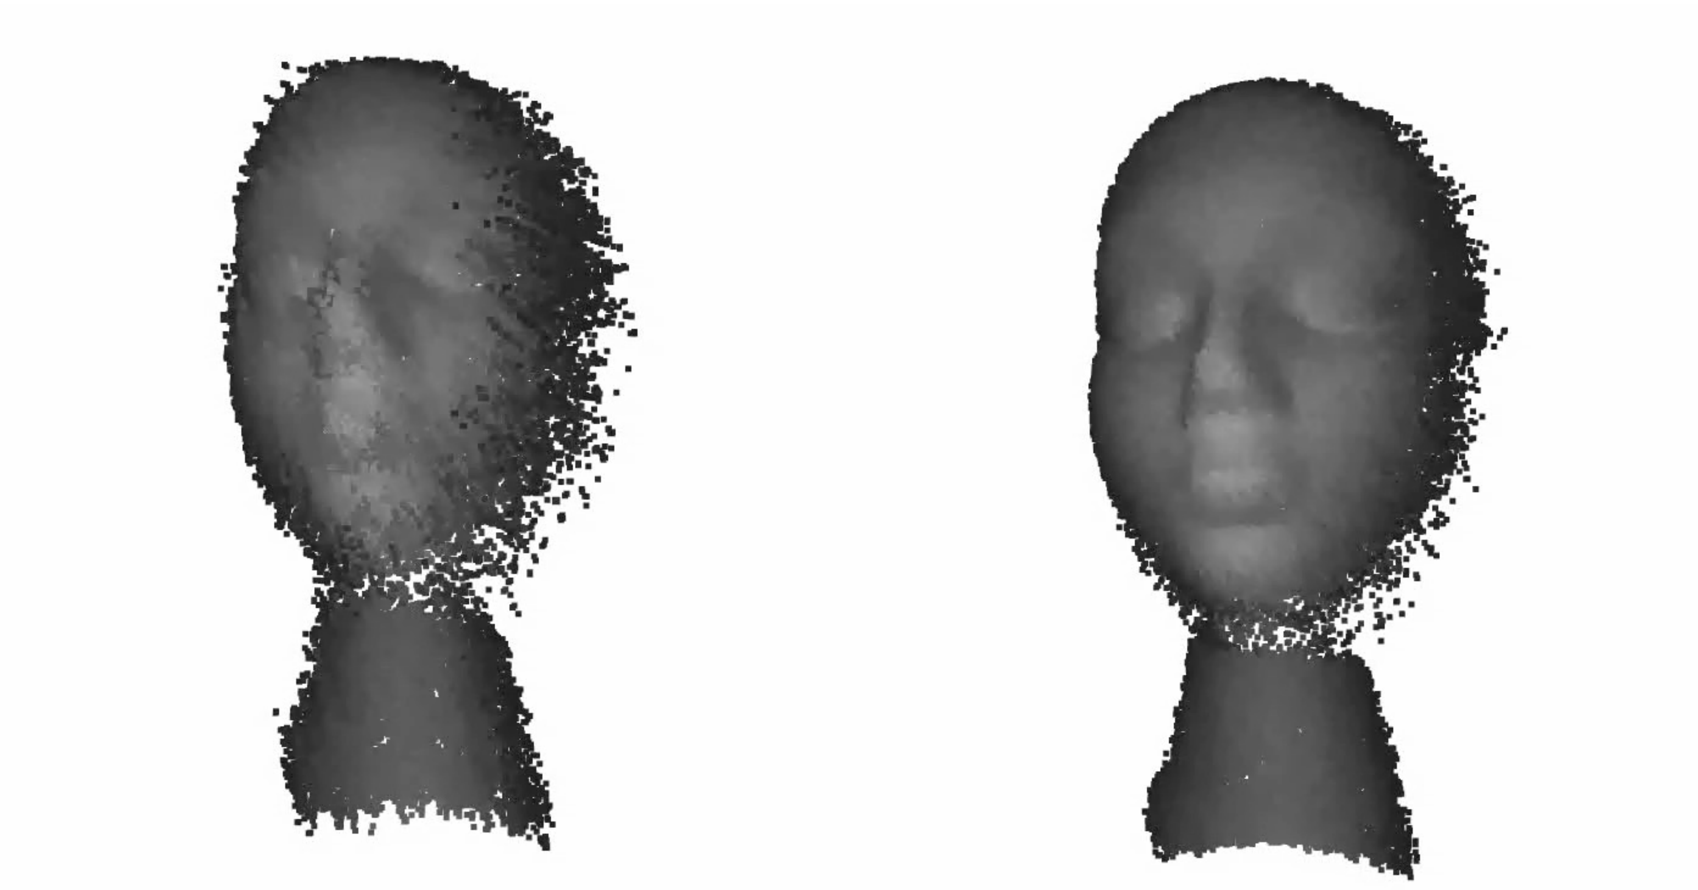

Low Power **This Work**

Real data captured using ADI ToF Camera*

*AD-96TOF1-EBZ

Summary of Contributions

- **Key Insight**: We can estimate the rigid motion using the sparse correspondences across the infrared images that a ToF camera collects
- Vary the amount of emitted light to infrequently obtain regular power depth maps and use them to denoise subsequent lower power ones
- Reduce the mean relative error of the low power depth maps by 64% in real-time on an embedded processor

Our Publication: ICIP 2019

Future Directions

- **Sensor Calibration:** Varying the amount of emitted light affects the temperature calibration

- **Saturation:** Objects that are close to the ToF camera can saturate the sensor, and subsequent depth maps may not be denoised

- **Other Noise Sources:** How do issues like multi-path interference affect this approach?

Outline

Strictly Computation

Strictly ToF Camera

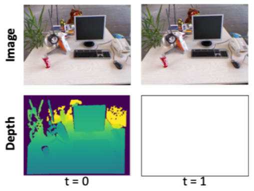

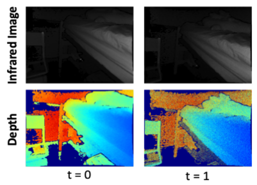

- Reduce the Usage of the ToF Camera
  - Depth Map Estimation for Rigid Scenes
  - Depth Map Estimation for Dynamic Scenes
- Reduce the Light the ToF Camera Emits
  - Adaptive Pulse Control
- **Summary of Thesis Contributions**

Summary of Thesis Conclusions

- By balancing the usage of the ToF camera with computation, we can lower the power required to obtain accurate and dense depth maps

- By exploiting rigidity, we can use sparse optical flow and linear least squares to reduce computation

- We show that our algorithms can estimate depth maps at up to real-time on embedded processors and that they can be used for real applications

Publications That Went Into Thesis

- **Conferences**
  - J. Noraky, V. Sze, "Low Power Depth Estimation for Time-of-Flight Imaging," ICIP, 2017.
  - J. Noraky, V. Sze, "Depth Estimation of Non-Rigid Objects for Time-of-Flight Imaging," ICIP, 2018.
  - J. Noraky, C. Mathy, A. Cheng, V. Sze, "Low Power Adaptive Time-of-Flight Imaging for Multiple Rigid Objects," ICIP, 2019.

- **Journal Publications and Preprints**
  - J. Noraky, V. Sze, "Low Power Depth Estimation of Rigid Objects for Time-of-Flight Imaging," TCSVT, 2020.
  - J. Noraky, V. Sze, "Depth Map Estimation of Dynamic Scenes Using Prior Depth Information," Under Review, 2020.
  - J. Noraky, V. Sze, “Low Power Depth Map Denoising for Mobile Time-of-Flight Cameras,” In Preparation, 2020.

Acknowledgements

- **Thesis Committee:** Vivienne Sze, Berthold K.P. Horn, Charles Mathy

- **Analog Devices**
  - Thank you for your generous funding and support!
  - Patrick Coady, Sefa Demirtas, Nicolas Le Dortz, Charles Mathy, Kaushal Sanghai, Daniel Schmidt, Jonathan Yedidia, Atulya Yellepeddi, Tao Yu

- **MIT Faculty and Staff:** Janice Balzer, Thomas Heldt, Gim Hom, Al Oppenheim, Terry Orlando, George Verghese

Acknowledgements

- **EEMS Group:** Alan, Amr, Diana, Gladynel, Hsin-Yu, Mehul, Nellie, Peter, Robert, Soumya, Tien-Ju, Theia, Yi-Lun, Yu-Hsin, Zhengdong

- **MIT Grad School**: Albert, Ali, Anne, Anuran, Ashwin, Austin, Curtis, Chai, Danielle, David, Eren, Ganesh, HanByul, Haoyang, Hoon, Isaak, Lehka, Lucas, Mary, Marek, Mohammed, Nirav, Orhan, Tuhin, Victor…

- **Family:** My mother, sister, and grandparents