

One Object at a Time: Accurate and Robust Structure From Motion for Robots

Aravind Battaje^{1,2} and Oliver Brock^{1,2}

Abstract—A gaze-fixating robot perceives distance to the fixated object and relative positions of surrounding objects immediately, accurately, and robustly. We show how fixation, which is the act of looking at one object while moving, exploits regularities in the geometry of 3D space to obtain this information. These regularities introduce rotation-translation couplings that are not commonly used in structure from motion. To validate, we use a Franka Emika Robot with an RGB camera. We a) find that error in distance estimate is less than 5 mm at a distance of 15 cm, and b) show how relative position can be used to find obstacles under challenging scenarios. We combine accurate distance estimates and obstacle information into a reactive robot behavior that is able to pick up objects of unknown size, while impeded by unforeseen obstacles.

I. INTRODUCTION

When humans move about, they fixate their gaze on objects of interest without conscious notice. Some theories suggest it may be to stabilize images on the retina, or expose those objects to the area of retina with highest visual acuity (fovea). We show that this action of fixating on one object at a time simplifies perception for any moving agent. It gives *robust* and *instantaneous* information crucial to interact with the fixated object: how far is it and how other objects are positioned with respect to it. This information can be easily adapted to various robotic behaviors, including picking objects in challenging scenarios such as in Fig. 1.

The concept and method we describe in the paper apply to actively moving agents. So, we use the term ‘fixation’ to mean fixating upon a point on an object *while* moving.

Fixation gives a typical robot access to relevant information about interacting with the world. It provides a way to measure **accurate distance** to an object by only observing robot’s self-movement, and robustly extract **relative positions** of surrounding objects by observing direction of visual motion on those objects.

These properties are also efficient to extract. A robot measures its own movement to control its position which directly helps calculate distances, and simple image tracking can sense direction of visual motion, which helps find relative positions of other objects. Moreover, a robot only needs to modify the way it moves: fixating on one object at a time, instead of moving in an unconstrained manner. Consequently, this movement permits it to act robustly in the world without needing a world model.

¹ Robotics and Biology Laboratory, Technische Universität Berlin

² Science of Intelligence, Research Cluster of Excellence, Berlin

We gratefully acknowledge funding provided by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2002/1 “Science of Intelligence” – project number 390523135.

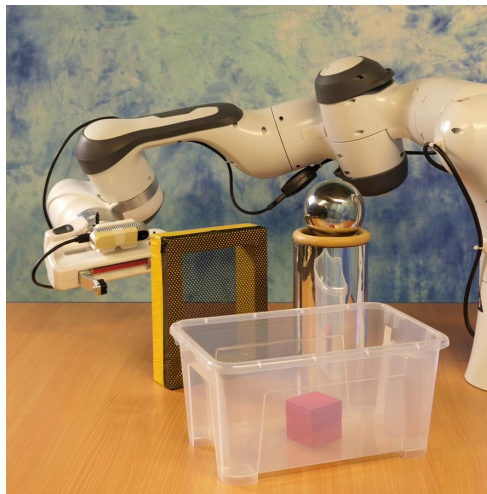


Fig. 1: Problems for vision in robotics: reflective and translucent objects; we present an approach that robustly avoids these objects while picking up an object of unknown size, using only an RGB camera

In this paper, we show how distance estimation using fixation outperforms that of a translating camera, and how determining obstacles is immune to common problems for vision such as translucent and reflective surfaces. We integrate these accurate and robust properties enabled by fixation into a real-world robotic application. We task a robot to pick up an object of unknown size impeded by unforeseen and dynamic obstacles. The robot can be seen in action¹ in Fig. 5.

II. RELATED WORK

In our method, we actively control a moving robot’s view in order to maintain fixation on one object at a time. By doing this, we extract relevant 3D properties of the world. So, our work relates to active vision, structure from motion (SfM) and other approaches that use camera fixations. In the following, we relate to these topics.

A. Structure From Motion

Our work lies in the intersection of SfM and active vision [1]. Recovering 3D properties from a moving camera is well-posed and stable [1] within active vision. But due to issues we explain below, depth recovery might not be accurate. Fixations can be seen as introducing constraints to overcome those issues.

¹A video is available at <https://youtu.be/o6xqkQINhmc>.

Typical approaches in SfM recover depth from the component of visual motion due to camera translations, as rotations alone do not reveal depth cues. Simpson says, “any process or mechanism for getting depth from optic flow must somehow filter out the rotational component” [2]. So, camera rotations are either removed [3] or compensated [4]. We show that camera rotations produced by fixation are indeed useful, especially in the context of robotics.

It is common in SfM to rely on point-features or “corners” [5], to circumvent correspondence issues due to the aperture problem [2] at edge features. But corners are also problematic: there might not be enough corner features on smooth objects, or too many on textured objects that might overwhelm corner matching. In contrast to using corners, we simply rely on points at edges that are easier to extract, and obtain closest point correspondences on them. This is sufficient to determine location of other objects through their direction of visual motion.

B. Fixation-Based Vision

Some works [6], [7], [8] have demonstrated superior depth recovery by reproducing characteristic fixation movement of humans on robots. Similar to our approach, they take advantage of the robustness achieved with fixation. In contrast, instead of building a depth map for the full field-of-view, we only retrieve information that is necessary for the robot to immediately take action, i.e., knowledge of impending objects and precise distance to the fixated object. We also show how these properties can be easily tailored to robust robot behaviors.

Fixation establishes a relationship between the scene and camera called the *Zero Flow Circle (ZFC)* [9], which can determine direction of heading [10]. We too use the ZFC but to filter objects “in front” of the fixated object from the rest of the visual scene.

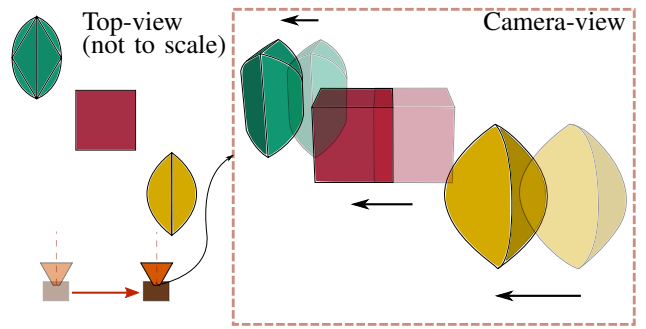
III. MOTION PARALLAX WITH GAZE FIXATION

In this section, we will look at the mechanics of gaze fixation, and how a camera moving in such a manner leads to robust properties.

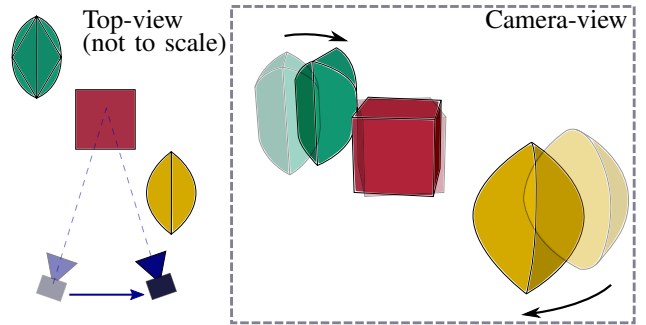
In order to understand fixation, let us first look at the effects of camera translation and rotation independently. Camera translation results in motion parallax as shown in Fig. 2a. Here, parallax is the apparent change in parts of the scene on the camera sensor surface, and distant objects produce lesser parallax. In contrast, camera rotations produce the same parallax for any object irrespective of distance.

Fixation combines camera translation and rotation in a unique way. In order to maintain an object in the same position in camera’s view, fixation produces a rotation that counteracts the effect of translation. In other words, the parallax due to rotation cancels that of translation. As a consequence, objects behind the fixated object produce parallax in the opposite direction as those in front, as seen in Fig. 2b.

This unique combination of translation and rotation produces robust information that is immediately helpful. The parallax in opposing directions serves as a filter to quickly



(a) Motion parallax with a translating camera



(b) Motion parallax with a moving camera fixating on the (center) red object. Due to fixation, the camera translates and rotates

Fig. 2: Translating camera leads to parallax only in one direction, whereas gaze fixation leads to parallax in different directions. Translucent colors indicate camera in starting position, and solid colors in the final position. (a) Farther objects move slower. Thus distance to an object can be found by measuring object translation for a given amount of camera translation. (b) Fixation produces motion in objects directly indicative of their position: fixated object moves minimally, objects behind this move in the opposite direction as objects in front. Distance to fixated object is the ratio of camera rotation to translation.

distinguish between objects that are in front, and behind the fixated object. The rotations are also dependent on distance to the object: farther objects induce lesser rotations to maintain fixation. Thus, the ratio between rotation and translation immediately reveals the **distance to the fixated object.**

We use these two principles as the basis to generate behaviors in a robot in Sec. IV. In this section, we explore these principles in more detail. The following two subsections explain the mathematical underpinnings, by first taking a look at the general case of camera motion, and later analyzing the influence of fixation.

A. Visual Motion Patterns Due to Camera Motion

Let us understand the general mathematical relation that expresses motion parallax. A camera projects 3D points onto the 2D sensor surface, and for a moving camera, parallax can be expressed as changes of those projected points for changes

of camera translation and rotation [2], [11].

The standard pinhole projection model relates a 3D point in the camera frame with coordinates $[X, Y, Z]^T$ to its 2D projection on the image plane $[x, y]^T$. This relation is given by the projection equations $x = f \cdot X/Z$ and $y = f \cdot Y/Z$. In rest of the paper, the focal length f is assumed to be unit length without loss of generality.

For the case of a moving camera, the velocity of projected points $[\dot{x}, \dot{y}]^T$ can be related to the camera velocity with the help of time derivative of projection equations. Eqn. 1 shows this relationship. The reader can find more information about this derivation in [11].

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = d \cdot \begin{bmatrix} -1 & 0 & x \\ 0 & -1 & y \end{bmatrix} \begin{bmatrix} v_X \\ v_Y \\ v_Z \end{bmatrix} + \begin{bmatrix} xy & -(1+x^2) & y \\ (1+y^2) & -xy & -x \end{bmatrix} \begin{bmatrix} \omega_X \\ \omega_Y \\ \omega_Z \end{bmatrix} \quad (1)$$

where $[v_X, v_Y, v_Z]^T$ is the camera translation velocity, $[\omega_X, \omega_Y, \omega_Z]^T$ is the camera rotation velocity, and $d = 1/Z$ is the disparity—the only quantity that relates to distance to 3D points. Since d influences the first term, often camera translation is regarded as helpful in extracting 3D properties, and rotations be filtered out [2]. On the contrary, we show that camera rotations are indeed useful.

B. Visual Motion Patterns Due to Gaze Fixation

Fixation can be seen as coupling two rotational degrees-of-freedom (DOF) with translations. So specifying camera velocities reduces to 4-DOF, instead of 6-DOF. This leads to a useful reduction of Eqn. 1 that results in patterns of parallax explained earlier. And the coupling relates to the distance of fixated object.

This coupling imposed by fixation is expressed as

$$\omega_X = -v_Y \cdot d_f \quad (2)$$

$$\omega_Y = v_X \cdot d_f \quad (3)$$

where $d_f = 1/Z_f$. This coupling entails camera rotation for any lateral translation to ensure a particular 3D point $[0, 0, Z_f]^T$ in the camera frame is held at the center of FOV $[0, 0]^T$. Please note, all quantities in this paper with subscript f denote their association with the *fixation point*. Therefore, Z_f is the distance to the fixated point.

For clear exposition, let us also set $v_Z = 0$ and $\omega_Z = 0$. The former leads to a quantity called time-to-collision (TTC) [2], which is not related to motion parallax. And the latter does not couple with fixation. In practice both can be accounted. Thus, setting irrelevant quantities to zero, and substituting Eqn. 2 and 3 in Eqn. 1 produces

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} (1+x^2)d_f - d & -xy d_f \\ xy d_f & (1+y^2)d_f - d \end{bmatrix} \begin{bmatrix} v_X \\ v_Y \end{bmatrix} \quad (4)$$

To understand the geometry induced by this expression, let us analyze the surface where projected velocities is zero, i.e., $[\dot{x}, \dot{y}]^T = 0$. This is a surface that contains the point of

fixation, but it also separates the space such that 3D points beyond this surface have the opposite sign of motion as those within the surface. This is also called the *Zero Flow Circle* [9].

For a robot interacting with the world, this separation is useful: based only on direction of motion, it can instantaneously determine parts of the scene “behind” the fixated point that can be simply ignored as far as direct interaction is concerned. Likewise, it can instantaneously determine parts of the scene that is “in front” of the fixated point that could perhaps be an obstacle.

In practice, we can make an approximation to Eqn. 4. If the camera has a small FOV, then the quadratic terms quickly vanish to zero: $x^2, y^2, xy \approx 0$. Consequently,

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \approx \begin{bmatrix} d_f - d & 0 \\ 0 & d_f - d \end{bmatrix} \begin{bmatrix} v_X \\ v_Y \end{bmatrix} \quad (5)$$

We use this approximation to construct a potential field for obstacles in Sec. IV-D by simply filtering points that have the opposite sign of camera movement (effectively $d_f < d$).

In the analysis so far, we introduced a new coupling (Eqn. 2, and 3), but it need not be pre-computed. Instead, the coupling comes about as a *result* of fixation. Fixation forces the camera to follow circles of different radius, depending on the distance to that 3D point of interest. The robot can then simply determine the instantaneous radius of curvature by examining its ratio of translation and angular velocities.

In summary, Eqn. 4 (or Eqn. 5)) describes the phenomenon of parallax of opposite directions we saw in the beginning of this section, and the coupling described in Eqn. 2 and 3 helps measure distance to fixated object.

IV. GAZE FIXATION FOR ROBUST ROBOT BEHAVIOR

Now let us look at the requirements for a robot to leverage the 3D properties enabled by fixation, and subsequently pick up objects of unknown size while avoiding obstacles. Later in this section, we explain our method that satisfies these requirements.

For purposes of simplicity, we will assume a single robot manipulator with a monocular camera sensor near its end-effector. The requirements we list below can be extended when camera can be moved independent of the manipulator.

The robot must

- 1) accomplish *fixation* by actively rotating in order to keep a target object in the center of FOV;
- 2) translate *laterally* to produce the special geometry induced by fixation;
- 3) *estimate distance* to target by measuring the ratio of translation and rotation;
- 4) observe *direction of movement* of objects relative to the target object;
- 5) *avoid obstacles* informed by the relative visual motion;
- 6) translate *towards* the target object to pick it up;
- 7) and, optionally, place the end-effector in a desirable configuration to facilitate successful grasp.

To turn these requirements into implementation, we use an instantiation of visual servoing [11] for 1, 2, 5, and 6.

This is not a generic visual servoing, but one that structures perceptual processes enabled by fixation 1. Then to enable reliable distance measurements 3, we integrate the robot state using an Extended Kalman Filter. And for 4, we use a simple form of optic flow. In the following subsections, we will explain our approach to visual servoing, optic flow, distance estimation and its use to avoid obstacles.

A. Behavior Generation Based on Gaze Control

We partition the 6-DOF control of a robot into three components as detailed below.

1) *Gaze fixation*: This maintains the location of a specified feature $[x_s, y_s]^T$ in the center of the image using the following servo law:

$$\omega_{XY} = \lambda_{XY} \begin{bmatrix} xy & -(1+x^2) \\ (1+y^2) & -xy \end{bmatrix}^{-1} e_{xy} \quad (6)$$

where $\omega_{XY} = [\omega_X, \omega_Y]^T$, and λ_{XY} is the proportional gain to minimize error $e_{xy} = [x_s, y_s]^T$. The feature $[x_s, y_s]^T$ may be any representative point within the area constituting an object. We use the centroid of the area occupied by a color-specified object as this feature². Finally, the Jacobian matrix is always invertible, which ensures stable fixation and robust fulfillment of requirement 1.

2) *Approach*: The approach servo law brings the robot to some specified distance in front of the fixated object:

$$v_Z = \lambda_Z \log \left(\frac{Z_d}{Z_f} \right) \quad (7)$$

where λ_Z specifies how fast to approach or retract from desired distance Z_d , and Z_f is the estimated fixation distance (estimation is explained in Sec. IV-C).

This servo law satisfies requirement 6 to move toward the object. To pick the target object up, we employ an open-loop maneuver: move the end-effector forward by Z_d and close the gripper.

3) *Translate laterally*: Lateral translations of the robot once an object is fixated serves two objectives: produce visual motion (requirement 2) and depending on existence of obstacles, move away from them (requirement 5). The following control law accomplishes both:

$$\mathbf{v}_{XY} = \lambda_{dist}(Z_f - Z_d) + \begin{bmatrix} \lambda_{cycle_X} \sin(t) \\ \lambda_{cycle_Y} \cos(t) \end{bmatrix} - \lambda_{obs} \mathbf{obs}_{xy} \quad (8)$$

where $\mathbf{v}_{XY} = [v_X, v_Y]^T$, and the three terms in control law are independent heuristics, weighted by their respective λ gains. The first term moves the robot laterally if it is far from the target, or if estimated distance to target is uncertain. The second term introduces small periodic movements. These two terms ensure motion needed to constantly determine distance to target object, and relative location of other objects. The last term moves the robot away from obstacles (more details in Sec. IV-D).

²This simple scheme also provides a stable fixation point for hollow objects.

The above three components are reactive control laws that hold the target object as temporary reference even when it moves, or obstacles change positions. Thus they result in a rich and robust behavior³.

B. Optic flow (requirement 4)

Optic flow algorithms enable tracking visual motion of points on objects. This helps determine position of objects relative to fixated object. However, since only direction of movement is important, the optic flow algorithm can be simple. We use optic flow derived from spatiotemporal gradients [12], [13].

C. Online Distance Estimation (requirement 3)

We employ an Extended Kalman Filter (EKF) to maintain a temporally consistent estimate of fixation point and associate uncertainty to it. The EKF filters over camera velocity and inverse distance to fixation point d_f . Its motion model integrates velocity with constant acceleration, and implements the fixation constraint from Eqn. 2 and 3. The EKF's observation model integrates camera velocity (obtained from robot's velocity) along with n samples of optic flow on the target object. We randomly sample $n = 30$ points on edge features of the color-specified object. The randomness allows sustained estimation with drastic changes in perspective.

The EKF also recovers quickly if the target moves, or the robot is unable to maintain fixation. As fixation maintains near-zero optic flow on the target object, any large deviations from it indicates inconsistencies, thus increasing uncertainty of estimate.

D. Avoiding Obstacles (aiding requirement 5)

As the presence of other objects "in front" of fixated object becomes apparent immediately through its direction of movement, this can instruct the robot to produce lateral translations (Sec. IV-A.3) to move in a direction with least influence of those objects. To enable this, we construct a potential function populated with high values for any pixels that have optic flow in the opposite sign of robot lateral movement, i.e., $[\dot{x}, \dot{y}] \cdot [v_X, v_Y]^T < 0$. We then calculate the center-of-mass of this potential field as \mathbf{obs}_{xy} , and instruct the robot to move in the *opposite* direction of it.

V. EXPERIMENTS

A. Experimental Setup

Here we describe the platform and parameters we use for all the upcoming sections. We use a Franka Emika Robot with a RealSense D435 camera rigidly attached near the end-effector—we only use the RGB sensor on the D435. The robot is controlled using the method described in previous section, in realtime at 1000 Hz. The optic flow and distance estimation run at approximately 25 Hz. We run all the above on a compact PC with an Intel Core i5 CPU.

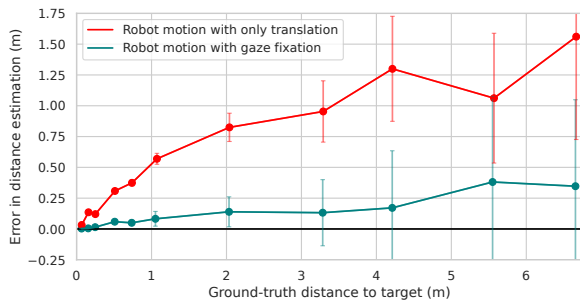
³Note, we leave out control of ω_Z . If desired, this maybe used in conjunction with an additional heuristic to Eqn. 8, to place the end-effector in some desirable configuration for pick up (requirement 7).

B. Gaze Fixation Estimates Distances Accurately

In this subsection, we evaluate and compare distance estimates from a fixating robot and a translating robot. This is similar to Fig. 2b and 2a respectively.



(a) Robot directed at a target in a motion capture system



(b) Distance estimation of a fixating robot outperforms that of a translating robot; Note, for visual clarity, we have flipped the sign of error for distance estimates under fixation

Fig. 3: Distance estimation experimental setup and results

The robot estimates distance to a target without moving toward it. Hence we disable approach (Sec. IV-A.2). The target is a printed “rubber duck”. These targets are printed at 5 different sizes to prevent the target from being too small ($< 5\%$ of image width) or too large ($> 75\%$ of image width) in the camera view at different distances. Our method has no information of the size of target. We vary the target distance two orders of magnitude, from 66 mm to 6.6 m. At each target distance, the robot end-effector translates laterally in a cyclic motion. Its amplitude is the same for both fixation and translation-only motion. To estimate distances, we use the EKF from Sec. IV-C. In the case of translation-only motion, we simply remove the *fixation constraint* in its motion model. To obtain the ground truth distances, the positions of robot end-effector and target are tracked in a high-precision motion capture system from OptiTrack, as shown in Fig. 3a. At each target distance, we report the mean and standard deviation of error between ground truth and estimated data from EKF, accumulated over approximately 9 seconds (225 samples).

Fig. 3b shows results of the evaluation. It shows that distance estimation error with fixation is significantly less than that when robot only translates. The error is less than 5 mm at around 15 cm distance—this would be a reasonable distance for a manipulator to start interacting with an object, e.g., pick it up. The error also grows slower with larger distances: it is less than 13 cm at around 2 m, which is only about 15% of translation-only motion. These impressive

results indicate that accurate distances using fixation can be utilized in real-world robots using only an RGB camera.

C. Gaze Fixation Is Robust in Challenging Scenarios

We will now examine how fixation leads to robust information about surrounding objects *while* estimating accurate distances to the fixated object. Fig. 4 shows the camera view and optic flow visualization for different scenarios. In each, the robot fixates on the red object. The opposing colors in optic flow visualization distinctly indicate relative positions of objects and scene surrounding the red fixated object—all objects in the front move in the opposite direction of those in the back. Even object reflections (fourth column) that appear to be in front can be clearly dismissed, because reflections move as if they were behind the fixated object. Finally, even though an obstacle might be physically in the line-of-sight of the fixated object (second column), distance to red object can be calculated independently as this only depends on the object itself being visible in camera view. As extracting this information requires only a basic optic flow implementation and no prior object models, this may be effortlessly integrated into robust robotic applications.

D. Gaze Fixation Produces Robust Robot Behavior

As we described in the previous subsections, using fixation, the robot can quickly extract robust information about surrounding objects, while estimating accurate distances to fixated object. We also described in Sec. IV an approach to combine this information and generate robot behavior. We demonstrate this on a robot¹ as shown in Fig. 5. Here the robot succeeds in picking up the target object, even when impeded by obstacles from Fig. 4, or if obstacles are moved mid-operation, or the target itself is moved. This leads us to conclude that fixation is a powerful behavior to incorporate in a robot, especially to build other robust behaviors on top of it.

VI. LIMITATIONS

Dependence on rotation: With fixation, distance recovery is dependent on amount of camera rotation: farther objects result in lesser rotations. Hence fidelity in measuring small camera rotation is paramount for farther objects. Without this, distance recovery will be erroneous. However, determining relative location of other objects does not impose restrictions on measuring camera rotations.

Stability of fixation: Unstable fixation point, e.g., due to noisy object detection, produces spurious camera rotations. Consequently distance estimation might not converge.

Ambiguity in fixation: We classify an object of interest using its color. This results in ambiguous fixation if there are other objects with the same color. This can be resolved with sophisticated object detectors and visual tracking.

Simplicity of obstacle avoidance: We employ a simple heuristic to avoid obstacles based on a potential function. This provides a powerful local motion primitive, but is prone to local minima. To avoid this, a motion planner could be used in conjunction to express even richer behaviors. We

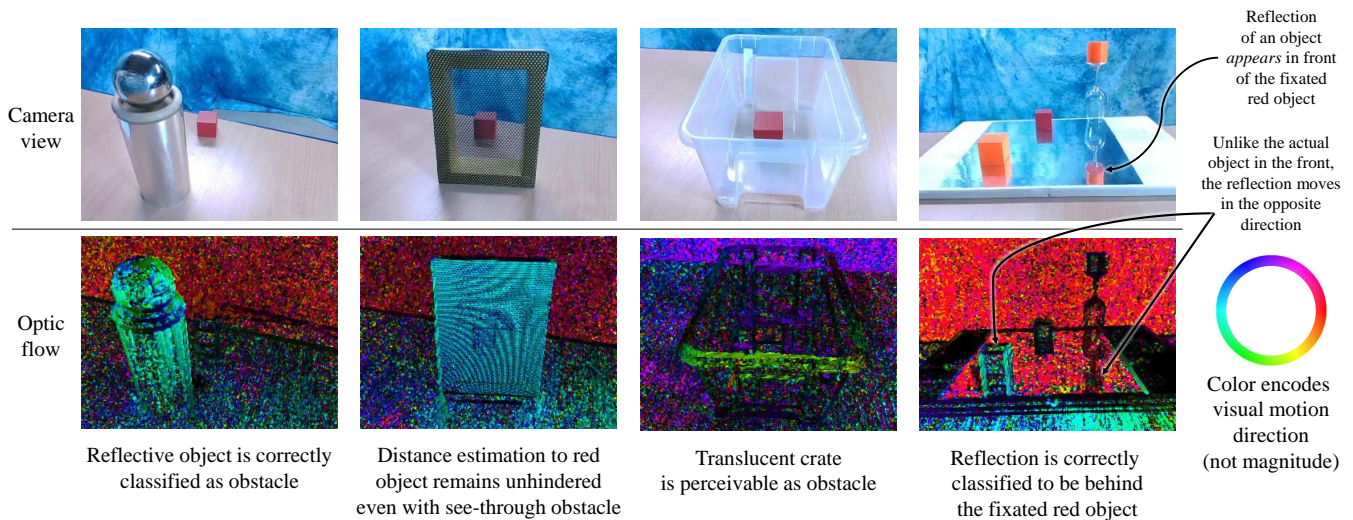


Fig. 4: Four scenarios difficult for vision but easy with fixation (more details in Sec. V-C)

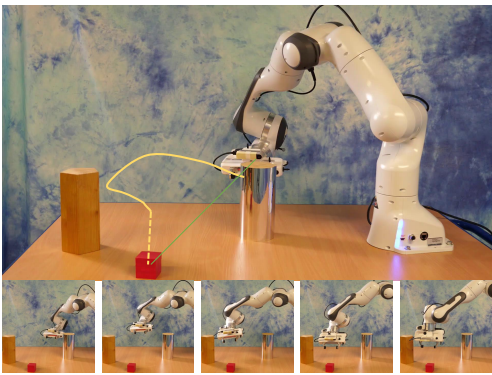


Fig. 5: Robot takes a complex path (yellow) to fetch the red object, in spite of several obstacles. The robot begins to avoid the reflective obstacle in its line-of-sight (green), and then avoids the wooden obstacle in the far left upon discovering it. Finally, the robot moves forward (dotted yellow) in an open loop fashion to grasp the red object.

have also ignored the contribution of TTC to this potential function, thus we maintain the approach of the robot relatively slower than lateral translations. TTC can be accounted as it results in a shifted ZFC [9] that still contains the fixation point.

VII. CONCLUSIONS

Gaze fixation leverages useful relationships between a moving agent and its surrounding world: it gives just-in-time access to distance to a fixated object, and relative location of other objects. We show how these properties emerge from fixation, and show that the resultant distance estimates are accurate and information about surrounding objects are robust. We also demonstrate a way to generate robot behavior utilizing these properties.

Fixation is also easy and efficient. It adds rotations that are not adversarial to a typical robot's objective; uses robot's

proprioception to measure distance; and utilize simple visual tracking to find relevant details of surrounding objects. It is also easy to adapt these properties into robust behaviors. We believe the perceptual and behavioral skills offered in this paper form a foundation for generating robust real-world robotic applications. We believe fixation has the potential to become ubiquitous in robotics.

REFERENCES

- [1] J. Aloimonos, I. Weiss, and A. Bandyopadhyay, "Active vision," *Intl. Journal of Computer Vision*, vol. 1, no. 4, pp. 333–356, Jan. 1988.
- [2] W. A. Simpson, "Optic flow and depth perception," *Spatial Vision*, vol. 7, no. 1, pp. 35–75, 1993.
- [3] P. K. Bideau, "Motion Segmentation - Segmentation of Independently Moving Objects in Video," *Doctoral Dissertations*, Mar. 2020.
- [4] S. Zingg, D. Scaramuzza, S. Weiss, and R. Siegwart, "MAV navigation through indoor corridors using optical flow," in *2010 IEEE Intl. Conf. on Rob. and Autom.*, May 2010, pp. 3361–3368, iSSN: 1050-4729.
- [5] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge: Cambridge University Press, 2004.
- [6] M. Antonelli, A. P. del Pobil, and M. Rucci, "Bayesian multimodal integration in a robot replicating human head and eye movements," in *2014 IEEE Intl. Conf. on Rob. and Autom. (ICRA)*. Hong Kong, China: IEEE, May 2014, pp. 2868–2873.
- [7] A. J. Duran and A. P. del Pobil, "Robot depth estimation inspired by fixational movements," *IEEE Trans. on Cognitive and Developmental Systems*, pp. 1–1, 2020.
- [8] F. Santini and M. Rucci, "Active estimation of distance in a robotic system that replicates human eye movement," *Robotics and Autonomous Systems*, vol. 55, no. 2, pp. 107–121, Feb. 2007.
- [9] D. Raviv, "A quantitative approach to camera fixation," in *1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Proceedings*, June 1991, pp. 386–392, iSSN: 1063-6919.
- [10] I. Thomas, E. Simoncelli, and R. Bajcsy, "Peripheral Visual Field, Fixation and Direction of Heading," in *Exploratory Vision*, R. C. Jain, M. S. Landy, L. T. Maloney, and M. Pavel, Eds. New York, NY: Springer New York, 1996, pp. 169–189, series Title: Springer Series in Perception Engineering.
- [11] F. Chaumette and S. Hutchinson, "Visual servo control, part I: Basic approaches," *IEEE Robotics & Automation Magazine*, vol. 13, no. 4, pp. 82–90, Dec. 2006.
- [12] E. P. Simoncelli, "Distributed representation and analysis of visual motion," Thesis, M.I.T., 1993, accepted: 2005-08-15T21:18:23Z.
- [13] E. Simoncelli, "Design of multi-dimensional derivative filters," in *Proceedings of 1st Intl. Conf. on Image Processing*, vol. 1, Nov. 1994, pp. 790–794 vol.1.