

Constant Space Complexity Environment Representation for Vision-based Navigation

Jeffrey Kane Johnson¹

Abstract—This paper presents a preliminary conceptual investigation into an environment representation that has constant space complexity with respect to the camera image space. This type of representation allows the planning algorithms of a mobile agent to bypass what are often complex and noisy transformations between camera image space and Euclidean space. The approach is to compute per-pixel potential values directly from processed camera data, which results in a discrete potential field that has constant space complexity with respect to the image plane. This can enable planning and control algorithms, whose complexity often depends on the size of the environment representation, to be defined with constant run-time. This type of approach can be particularly useful for platforms with strict resource constraints, such as embedded and real-time systems.

I. INTRODUCTION

A significant issue in planning and control when solving real-world navigation problems is that there are often large numbers of individual agents with whom a mobile robot might interact. Consider navigating a busy roadway or crowded sidewalk or convention hall, where there may be multitudes of other agents sharing the space. Conventional approaches to planning in multi-agent systems often explicitly consider interactions between all agents and so become overwhelmed as the number of agents grows [1], [2], [3]. More scalable conventional approaches often have strict requirements on system dynamics [4] or observability of agent policies [5].

This paper presents a preliminary conceptual investigation into the use of a fixed-size environment representation for vision-based navigation. The representation is modeled after a camera image space, which is chosen because cameras are a ubiquitous sensor modality, and image space is by nature discrete and fixed size. The proposed representation additionally allows planning and control routines to reason almost directly in sensor space thereby avoiding often complex and noisy transformations to and from a more conventional Euclidean space representation. The intent of this new representation is to help vision-based mobile robots navigate complex multi-agent systems efficiently, and to take a step toward satisfying the strict resource requirements often present in real-time, safety critical, and embedded systems [6].

The next section briefly surveys related work, then the environment representation is presented along with an illustrative example of how it can be used. Finally, conclusions and future work are discussed.

¹Jeffrey Kane Johnson received his PhD from Indiana University and is principal of Maeve Automation, Mountain View, CA 94043
contact@maeveautomation.com

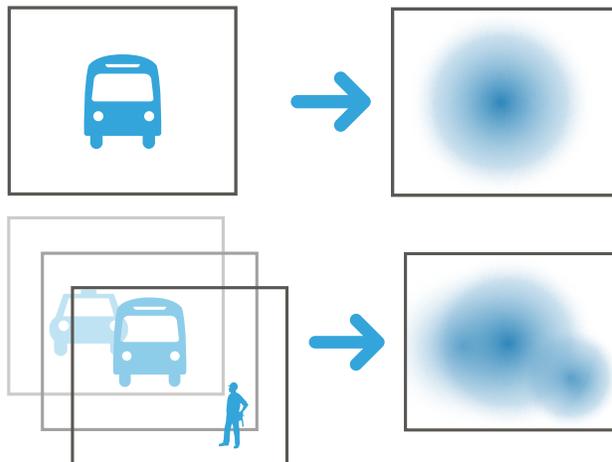


Fig. 1: Top: Illustration of an approaching object in the image plane (left), and the image space potential field (right). Bottom: Multiple object detections (left) can be composed into a single field (right). Black boxes represent ROIs.

II. RELATED WORK

The approach in this work is based on *potential fields* [7]. These fields represent attractive and repulsive forces as scalar fields over a robot’s environment that, at any given point, define a force acting upon the robot that can be interpreted as a control command. This type of approach is subject to local minima and the *narrow corridor* problem [8], particularly in complex, higher-dimensional spaces [9]. Randomized approaches can partially overcome these difficulties [10], [11], and extensions to globally defined *navigation functions* [12], [13], [14], while often difficult to use in practice, theoretically solve them. This work uses potential fields defined over a virtual image plane, which limits the possibility of narrow corridors, and is designed such that additional information can be used by the controller to interpret potential gradients, as suggested in [15]. [16] described a scheme related to that presented in this paper, but places potentials in a Euclidean space, which this approach explicitly avoids. The potential fields computed by the approach in this paper are intended to inform a *visual servoing* [17], [18], [19], [20] control scheme. This is a class of controllers that computes control commands directly from image space data.

In order to define values for the potential fields, this approach draws on a wealth of related works in optical flow and monocular collision avoidance, notably [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [24], [23]. The intuition of these approaches is that the information contained in a

sequence of monocular stills provides sufficient information to compute *time-to-contact* (Definition 2), which informs an agent about the rate of change of object proximity. The primary contribution of this work is a sensor-inspired representation space and algebra for enabling planning and control algorithms to reason effectively and efficiently with the output of this class of perception algorithms.

III. IMAGE SPACE POTENTIAL FIELDS

Before defining image space potential fields, the notion of a potential field used in this paper is defined below:

Definition 1. A *potential field* (also *artificial potential field*) is field of artificial forces that attracts toward desirable locations and repels from undesirable locations. In this work, a potential field is defined by a potential function that maps an image pixel value $I(x,y)$ to a tuple of affinely extended reals $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$, the first of which is the potential value, and the second of which is its time derivative:

$$I(x,y) \mapsto \overline{\mathbb{R}}^2 \quad (1)$$

From the definition, the image space potential (ISP) field is modeled after an image plane. As with image planes, the potential field is discretized into a grid, and regions of interest (ROIs) are defined for it. In this work it is assumed that the fields are origin and axis aligned with the camera images, and that they have the same ROIs (as in Figure 1).

The potential function, which maps image pixel values to potential values, can be defined in arbitrary ways, either with geometric relations, learned relations, or even heuristic methods. In this paper geometric properties of temporal image sequences are used. The approach is to assume an estimation of the *time-to-contact* (defined below) is available for each pixel in an image over time. The mapping of image pixel to this value is taken as the potential function that defines the image space potential field.

Definition 2. *Time-to-contact* (τ), is the predicted duration of time remaining before an object observed by a camera will come into contact with the image plane of the camera. The time derivative of τ is written $\dot{\tau}$.

As noted often in literature (e.g. [30], [25], [27], [29]), τ can be computed directly from the motion flow of a scene, which is defined as:

Definition 3. *Motion flow* is the pattern of motion in a scene due to relative motions between the scene and the camera. In other words, it is a vector field describing the motion of objects on the image plane over time.

Unfortunately, it is typically not possible to measure motion flow directly, so it is usually estimated via *optical flow*, which is defined as the *apparent* motion flow in an image plane. Historically this has been measured by performing some kind of matching of, or minimization of differences between, pixel intensity values in subsequent image frames [31], [28], [32], while more recently deep learning techniques have been successfully applied [33].

The image space potential field is now defined using τ :

Definition 4. An *image space potential field* is defined by a potential function that maps image pixels to a tuple of scalar potential values $\langle \tau, \dot{\tau} \rangle$.

A. Computing τ

Assuming some reasonably accurate estimation of optical flow vector field exists, τ can be computed directly under certain assumptions [25]. In practice, the computation of optical flow tends to be noisy and error prone, so feature- and segmentation-based approaches can be used [24], [23]. The idea of these approaches is to compute τ from the rate of change in detection *scale*. For a point in time, let s denote the scale (maximum extent) of an object in the image, and let \dot{s} be its time derivative. When the observed face of the object is roughly parallel to the image plane, and under the assumption of constant velocity translational motion and zero yaw or pitch, it is straightforward to show that [34]:

$$\tau = \frac{s}{\dot{s}} \quad (2)$$

As shown by Lemma 1, scale has a useful invariance property for these types of calculations that can make τ computations robust to certain types of detection noise:

Lemma 1. *The scale s of an object on the image plane is invariant to transformations of the object under $SE(2)$ on the XY plane.*

Proof. Let (X_1, Y_1, Z) and (X_2, Y_2, Z) be end points of a line segment on the XY plane in the world space, with XY parallel to the image plane and Z coincident with the camera view axis. Without loss of generality, assume unit focal length. The instantaneous scale s of the line segment in the image plane is given by:

$$s = \frac{1}{Z} \sqrt{\Delta X^2 + \Delta Y^2} \quad (3)$$

Thus, any transformation of the line segment on the XY plane for which $\Delta X^2 + \Delta Y^2$ is constant makes s , and thereby \dot{s} and τ , independent of the values of (X_1, Y_1) and (X_2, Y_2) . By definition, this set of transformations is $SE(2)$. \square

In addition, as shown in [35], the time derivative $\dot{\tau}$ of τ , when available, enables a convenient decision function for whether an agent's current rate of deceleration is adequate to avoid head-on collision or not. The decision function is given below, where $\epsilon > 0$ is a buffer to prevent an agent from coming to a stop directly at the point of contact with another agent:

$$f(\dot{\tau}, \epsilon) = \begin{cases} 1 & : \dot{\tau} \geq -0.5 + \epsilon \\ 0 & : \dot{\tau} < -0.5 + \epsilon \end{cases} \quad (4)$$

Equation 2 allows the computation τ for whole regions of the image plane at once given a time sequence of labeled image segmentations, and Equation 4 enables binary decisions to be made about the safeness of the agent's current

state. The following two sections describe encoding these pieces of information into image space potential fields.

B. Computing Fields for a Single Object

Computing the image space potential field for a single object is straightforward given the discussion in §III-A. Assuming an object can be accurately tracked and segmented over time in the camera image frame, its scale s and estimated expansion \dot{s} can be used to compute τ for each pixel in the image plane that belongs to the object, and a finite differences or estimation method can be used to compute $\dot{\tau}$. Pixels that do not belong to the object, and for which no other information is available, are mapped to $\langle \infty, \infty \rangle$ by the potential function. An illustration of this mapping is shown in the top row of Figure 1.

C. Computing Fields for Arbitrary Objects

Computing the image space potential field for arbitrary objects builds on the single object case by computing the field individually for each segmented and tracked object and then composing them into a single field. For this composition to be meaningful, however, the fields cannot be simply added together; this would result in the destruction of the τ information. Instead, a composite field is defined to preserve and combine τ information meaningfully. Equation 5 defines a composite field F in terms of image space potential fields F_1 and F_2 for an image I , and where \min_{τ} selects the tuple whose τ value is minimum:

$$F(x,y) = \left\{ \min_{\tau} (F_1(x,y), F_2(x,y)) \mid (x,y) \in I \right\} \quad (5)$$

Selecting the point-wise τ -minimum tuple for the composite field effectively enforces instantaneous temporal ordinality of objects, i.e., objects that are instantaneously temporally nearer are always what is seen. It is important to note that this is **not** the same as spatial ordinality. For an illustration of this, see Figure 2.

D. Constant Space Complexity

By definition the image space potential field representation has guaranteed constant space complexity assuming that the camera images for which the fields are generated are fixed size. This can be a powerful tool in simplifying planning and control algorithms whose complexity is typically dependent on the number of objects in a scene. In many cases it may, in fact, be possible to achieve constant time for planning and control given this representation.

It is important to note, however, that computing the representation itself may have arbitrary complexity: the problem of segmenting and tracking objects in order to generate these potential fields, for instance, can be efficient or arbitrarily complex, depending on the approach. The problem of investigating efficient computation methods for these fields is a point of future work discussed in §V.

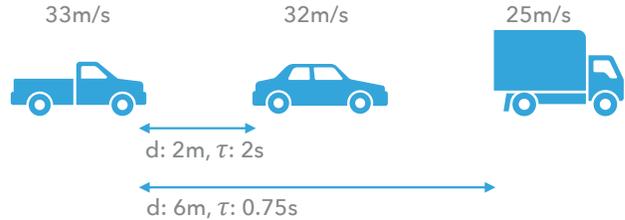


Fig. 2: Illustration comparing spatial and temporal ordinality. Consider three vehicles traveling in the same lane. For the pickup truck (left), the car (middle) has lowest spatial ordinality, i.e., is closest spatially. However, the van (right) has lowest temporal ordinality, i.e., it is nearest temporally.

IV. NAVIGATION WITH IMAGE SPACE POTENTIAL FIELDS

The intuition behind using image space potential functions for vision-based navigation is that they provide a natural space in which to compute collision avoiding controls, which then allows general navigation to be solved using a guided collision avoidance scheme, such as ORCA [4] or the Selective Determinism framework [36]. This section uses the Selective Determinism framework to define a simple control function that utilizes image space potential fields to navigate toward a visible goal in the presence of other agents. In this example only forward field of view is considered, but the extension to omnidirectional field of view is straightforward. The navigation problem considered is defined below:

Problem 1. Assume a set of acceleration-bounded agents \mathcal{A} , each operating according to known dynamics and with similar capabilities, navigating a shared space. Each agent operates according to a unique reward function that is not observable to other agents. Each agent is equipped with a camera that always faces in the direction of motion, and each agent is capable of performing image segmentation on the output. Suppose the reward function for an agent A encourages navigating toward a goal that A has in line of sight. How can A control toward the goal while avoiding collision with other agents?

Problem 1 is the type of problem that a driver may face on a busy highway when trying to navigate toward an exit or offramp. The solution in this example will take a naïve approach of decoupled steering and acceleration control while noting that more sophisticated control schemes are certainly possible. And while the example is formulated for a mobile agent traveling on a two dimensional manifold, the technique in general is equally applicable to three dimensions (such as with a UAV). The method for computing collision avoiding controls is discussed first, followed by the formulation of the navigation law.

A. Collision Avoidance

In order to address collision avoidance, the *Encroachment Detection* problem is presented.

Problem 2. Let *encroachment* refer to the reduction in minimum proximity between two or more objects in a workspace \mathcal{W} beyond desired limits as measured by a metric $\mu(\cdot, \cdot)$. Assume an agent A receives some observation input O_t of \mathcal{W} over time. Let \mathcal{A} be the set of agents that does not include A . For a sequence of observations O_i, \dots, O_t , how can A estimate the rate of change of $\min_{A_j \in \mathcal{A}} \mu(A, A_j)$?

Note that maintaining an estimate of $\langle \tau, \dot{\tau} \rangle$ directly solves the problem, as these values quantify temporal proximity and the rate of encroachment. The collision avoidance problem can now be solved by detecting encroachment and controlling such that it does not violate limits.

It was shown in [37] that collision avoidance can be guaranteed in a non-adversarial system if all agents compute and maintain collision-free stopping paths, which are contingency trajectories that bring an agent to zero relative velocity in minimal time. If agents are also self-preserving, they can each assume that all other agents will maintain such contingencies. Under these assumptions, agents should have sufficient information in the image space potential field to compute a solution to Problem 2 by maintaining non-zero time headway, which is assumed to be witness to the existence of a feasible stopping path.

For illustration, a naïve control set computation in the spirit of the braking controller in [38], [39] is sketched in Algorithm 1. This routine makes the reasonable assumption that $\dot{\tau}$ is not so large as to overwhelm τ . The idea is that steering angle and acceleration commands are computed independently and such that τ thresholds are not violated. To compute steering angles, a min filter is swept across the field of view in the potential field and a minimum potential value within the window is computed for each column in the image. Any value that meets τ thresholds is kept, and these are considered the safe steering angles (Figure 3, left). To compute the acceleration command, the minimum potential value within a centered window of a specified width is considered (Figure 3, right). If the value meets the τ threshold, the full scaled range of accelerations, $[-1, 1]$, is considered safe. If the threshold is violated, then either full deceleration $[-1, -1]$ or the range of deceleration values $[-1, 0]$ is sent depending on the value of the decision function of Equation 4. The control sets are then used by the Selective Determinism framework to compute the output control command.

B. The Selective Determinism Framework

Selective Determinism [36] is a solution framework for dynamically-constrained, non-adversarial, partially-observable multi-agent navigation problems. It belongs to a family of approaches useful for dealing with real-world problems because they remove the theoretical intractability inherent in optimal approaches [40], [41] while typically exhibiting good empirical performance. Selective Determinism, in addition, can also make certain collision avoidance guarantees even without explicitly considering interaction effects among agents.

Algorithm 1 Given an image space potential field F , compute the set of steering and acceleration commands that satisfy $\tau \geq T_s$ and $\dot{\tau} \geq -0.5 + \epsilon$, where $T_s > 0$ is some desired time headway, w_θ and w_a are kernel widths for computing steering angle and acceleration maps, and $\epsilon > 0$ is the buffer from Equation 4.

```

1: procedure SAFECONTROLS( $F, T_s, \dot{\tau}_E, w_\theta, w_a, \epsilon$ )
2:   Let  $I_c$  be the list of image column indices
3:   Let  $M_a$  map  $i \in I_c$  to steering angles
4:   Let  $h$  be the height (row count) of  $F$ 
5:   Let  $M_\tau$  map  $\langle \tau, \dot{\tau} \rangle$  to  $i \in I_c$  via  $w_\theta \times h$  min filter
6:   Let  $M_\theta = \{ \langle \tau, \dot{\tau} \rangle \in M_\tau : \tau \geq T_s \}$ 
7:   Let  $W$  be a centered  $w_a \times h$  window in  $F$ 
8:   Let  $\langle \tau, \dot{\tau} \rangle_{\min}$  be the min.  $\tau$  over  $W$ 
9:   Let  $L \leftarrow \emptyset$  be a container for safe accelerations
10:  if  $M_\theta = \emptyset$  then
11:     $M_\theta \leftarrow 0, L \leftarrow [-1, -1]$ 
12:  else if  $\tau_{\min} > T_s$  then
13:     $L \leftarrow [-1, 1]$ 
14:  else
15:    if  $f(\dot{\tau}, \epsilon) = 0$  then
16:       $L \leftarrow [-1, -1]$ 
17:    else
18:       $L \leftarrow [-1, 0]$ 
19:    end if
20:  end if
21:  return  $M_\theta, L$ 
22: end procedure

```

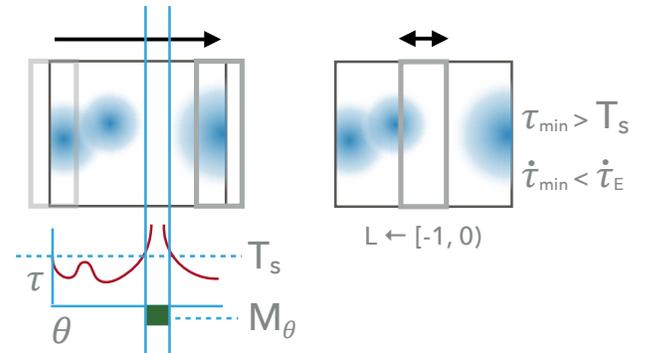


Fig. 3: Illustration of the steering angle control computation (left) and the acceleration control computation (right). On the left, a window sweeps from left to right over the image space potential field computing minimum τ for each image space column (left bottom). The set of column values that satisfy the threshold are the set of acceptable steering angles M_θ . On the right, the minimum potential value over a centered window is computed and the set L of acceptable scaled acceleration values are determined from it.

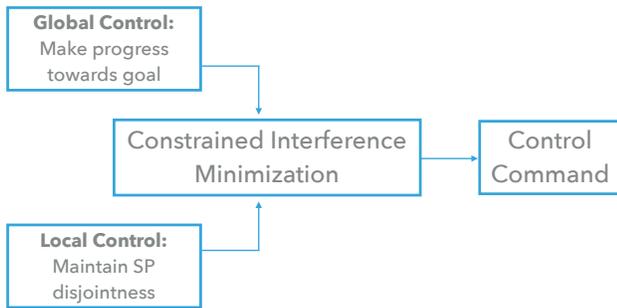


Fig. 4: Architecture of the Selective Determinism framework.

Selective Determinism works by exploiting the idea that agents in a system are capable of independently computing contingency trajectories that cover the space necessary for them to come to a stop (or to a zero relative velocity), and it assumes that agents do so, and that they will seek to maintain a non-empty set of available contingencies.

The framework casts the navigation problem in terms of a constrained interference minimization [38], [39] that utilizes a local collision avoidance controller to compute sets of controls from which an optimization chooses a control that makes maximal progress toward some goal (see Figure 4). The solution to Problem 1 is sketched in Algorithm 2.

Algorithm 2 For a desired pixel location (x_d, y_d) , and setpoint speed \dot{s}_d , compute the Selective Determinism control that safely guides the agent A toward (x_d, y_d) . See Algorithm 1 for descriptions of the other parameters.

```

1: procedure CONTROLS( $(x_d, y_d), F, T_s, \dot{\tau}_E, w_\theta, w_a, \epsilon$ )
2:   Let  $\theta_t, \dot{s}_t$  be the steering angle and speed of  $A$ 
3:   Let  $\theta_d$  be the steering angle corresponding to  $y_d$ 
4:   Let  $M_\theta, L \leftarrow \text{SafeControls}(F, T_s, \dot{\tau}_E, w_\theta, w_a, \epsilon)$ 
5:   Let  $\theta^* \leftarrow \theta_t$  contain the new steering angle
6:   for  $\theta \in M_\theta$  do
7:     if  $|\theta - \theta_d| < |\theta^* - \theta_d|$  then
8:        $\theta^* \leftarrow \theta$ 
9:     end if
10:  end for
11:  Let  $\dot{s}^* \in L$  be chosen proportionally to  $\dot{s}_d - \dot{s}_t$ 
12:  return  $\theta^*, \dot{s}^*$ 
13: end procedure
  
```

C. Complexity Analysis

In Algorithm 1 all non-trivial operations are iterations over the width of the image plane, which is assumed to be fixed for a given problem. The operations on lines 5 & 7 depend on the user defined w_θ and w_a parameters, but these are also bounded by image width. In Algorithm 2, Line 4 is a call to Algorithm 1, and so has constant complexity with respect to the image space, and Line 11 is assumed to be implemented with an $O(C)$ proportional law. Thus, the navigation algorithm as a whole has constant complexity, in space and time, with respect to the camera image space.

V. CONCLUSION & FUTURE WORK

This paper presented a conceptual investigation into an environment representation for vision-based navigation that has constant space complexity with respect to the image. This preliminary work is intended to serve as a basis for future investigations. This section outlines three primary topics of investigation.

The first topic is how to more completely combine environment information with the potential fields. As presented here, the representation is defined strictly in terms of object τ values, but a more elegant solution would build richer information about the environment into the potential field itself. An obvious extension would be to encode path information, such as lane or sidewalk boundaries, as well as goal information, into the potential field. This would enable navigation in semantically sophisticated environments and is an area of active development¹.

The second topic is a more sophisticated control law. The decoupled approach used here can lead to odd and counter-productive behavior, such as swerving out of the way of approaching objects while at the same time slowing down, or instability around control modes. A more intelligent control law would address stability issues and reason about the steering and longitudinal controls simultaneously. Additionally, allowing the potential fields to label a small, fixed-size set of objects individually could let such a control law reason about individual interactions without losing constant space complexity or information about all other objects.

Finally, the third topic, and one of great importance, is whether and how the potential fields themselves can be computed with some kind of constant complexity. A purely optical flow based approach would address this, but would require breakthroughs in the quality and efficiency of optical flow algorithms. Alternatively, a purely learning-based approach in conjunction with cheap, heuristic-based tracking approaches may provide the requisite segmentation and tracking information without runaway complexity.

REFERENCES

- [1] S. Brechtel, T. Gindele, and R. Dillmann, "Probabilistic mdp-behavior planning for cars," in *14th International IEEE Conference on Intelligent Transportation Systems, ITSC 2011, Washington, DC, USA, October 5-7, 2011*, 2011, pp. 1537–1542. [Online]. Available: <https://doi.org/10.1109/ITSC.2011.6082928>
- [2] E. Galceran, A. G. Cunningham, R. M. Eustice, and E. Olson, "Multipolicy decision-making for autonomous driving via changepoint-based behavior prediction: Theory and experiment," *Auton. Robots*, vol. 41, no. 6, pp. 1367–1382, 2017. [Online]. Available: <https://doi.org/10.1007/s10514-017-9619-z>
- [3] F. A. Oliehoek and C. Amato, *A Concise Introduction to Decentralized POMDPs*, ser. Springer Briefs in Intelligent Systems. Springer, 2016.
- [4] J. van den Berg, S. J. Guy, M. C. Lin, and D. Manocha, "Reciprocal n -body collision avoidance," in *Robotics Research - The 14th International Symposium, ISRR, August 31 - September 3, 2009, Lucerne, Switzerland, 2009*, pp. 3–19.
- [5] K. E. Bekris, D. K. Grady, M. Moll, and L. E. Kavraki, "Safe distributed motion coordination for second-order systems with different planning cycles," *I. J. Robotic Res.*, vol. 31, no. 2, pp. 129–150, 2012. [Online]. Available: <http://dx.doi.org/10.1177/0278364911430420>

¹https://bitbucket.org/maeveautomation/maeve_automation.core/src

- [6] S. Louise, V. David, J. Delcoigne, and C. Aussaguès, "OASIS project: deterministic real-time for safety critical embedded systems," in *Proceedings of the 10th ACM SIGOPS European Workshop, Saint-Emilion, France, July 1, 2002*, 2002, pp. 223–226. [Online]. Available: <http://doi.acm.org/10.1145/1133373.1133419>
- [7] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation, St. Louis, Missouri, USA, March 25-28, 1985*, pp. 500–505. [Online]. Available: <https://doi.org/10.1109/ROBOT.1985.1087247>
- [8] Y. Koren and J. Borenstein, "Potential field methods and their inherent limitations for mobile robot navigation," in *IN PROC. IEEE INT. CONF. ROBOTICS AND AUTOMATION*, 1991, pp. 1398–1404.
- [9] S. M. LaValle, *Planning algorithms*. Cambridge University Press, 2006.
- [10] *A Monte-Carlo algorithm for path planning with many degrees of freedom*, 1990. [Online]. Available: <http://ieeexplore.ieee.org/xpls/abs.all.jsp?arnumber=126256>
- [11] J. Barraquand and E. Latombe, "Robot motion planning: A distributed representation approach," *I. J. Robotics Res.*, vol. 10, no. 6, pp. 628–649, 1991. [Online]. Available: <https://doi.org/10.1177/027836499101000604>
- [12] D. Koditschek and E. Rimon, "Robot navigation functions on manifolds with boundary," *Advances in Applied Mathematics*, vol. 11, no. 4, pp. 412–442, 1990.
- [13] D. C. Conner, A. A. Rizzi, and H. Choset, "Composition of local potential functions for global robot control and navigation," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, vol. 4, Oct 2003, pp. 3546–3551 vol.3.
- [14] D. E. Koditschek, "Exact robot navigation by means of potential functions: Some topological considerations," in *Proceedings of the 1987 IEEE International Conference on Robotics and Automation, Raleigh, North Carolina, USA, March 31 - April 3, 1987*, pp. 1–6. [Online]. Available: <https://doi.org/10.1109/ROBOT.1987.1088038>
- [15] A. A. Masoud, "Solving the narrow corridor problem in potential field-guided autonomous robots," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation, ICRA 2005, April 18-22, 2005, Barcelona, Spain, 2005*, pp. 2909–2914. [Online]. Available: <https://doi.org/10.1109/ROBOT.2005.1570555>
- [16] M. T. Wolf and J. W. Burdick, "Artificial potential functions for highway driving with collision avoidance," in *2008 IEEE International Conference on Robotics and Automation, ICRA 2008, May 19-23, 2008, Pasadena, California, USA, 2008*, pp. 3731–3736. [Online]. Available: <https://doi.org/10.1109/ROBOT.2008.4543783>
- [17] N. J. Cowan and D. E. Chang, "Toward geometric visual servoing," EECB Department, University of California, Berkeley, Tech. Rep. UCB/CSD-02-1200, Sep 2002. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2002/5822.html>
- [18] L. E. Weiss, A. C. Sanderson, and C. P. Neuman, "Dynamic visual servo control of robots: An adaptive image-based approach," in *Proceedings of the 1985 IEEE International Conference on Robotics and Automation, St. Louis, Missouri, USA, March 25-28, 1985*, pp. 662–668. [Online]. Available: <https://doi.org/10.1109/ROBOT.1985.1087296>
- [19] R. T. Fomena and F. Chaumette, "Visual servoing from spheres using a spherical projection model," in *2007 IEEE International Conference on Robotics and Automation, ICRA 2007, 10-14 April 2007, Roma, Italy, 2007*, pp. 2080–2085. [Online]. Available: <https://doi.org/10.1109/ROBOT.2007.363628>
- [20] A. X. Lee, S. Levine, and P. Abbeel, "Learning visual servoing with deep features and fitted q-iteration," *CoRR*, vol. abs/1703.11000, 2017. [Online]. Available: <http://arxiv.org/abs/1703.11000>
- [21] A. Schaub and D. Burschka, "Reactive avoidance of dynamic obstacles through optimization of their epipoles," in *2015 19th International Conference on System Theory, Control and Computing (ICSTCC)*, Oct 2015, pp. 318–324.
- [22] G. Alenyà, A. Nègre, and J. L. Crowley, "Time to contact for obstacle avoidance," in *Proceedings of the 4th European Conference on Mobile Robots, ECMR'09, September 23-25, 2009, Mlini/Dubrovnik, Croatia, 2009*, pp. 19–24.
- [23] J. Byrne and C. J. Taylor, "Expansion segmentation for visual collision detection and estimation," in *IEEE International Conference on Robotics and Automation, Kobe, Japan, May 12-17, 2009*, pp. 875–882.
- [24] T. Mori and S. Scherer, "First results in detecting and avoiding frontal obstacles from a monocular camera for micro unmanned aerial vehicles," in *2013 IEEE International Conference on Robotics and Automation, Karlsruhe, Germany, May 6-10, 2013*, pp. 1750–1757.
- [25] T. Camus, D. Coombs, M. Herman, and T. Hong, "Real-time single-workstation obstacle avoidance using only wide-field flow divergence," in *13th International Conference on Pattern Recognition, ICPR 1996, Vienna, Austria, 25-19 August, 1996*, 1996, pp. 323–330.
- [26] S. J. Pundlik, E. Peli, and G. Luo, "Time to collision and collision risk estimation from local scale and motion," in *Advances in Visual Computing - 7th International Symposium, ISVC 2011, Las Vegas, NV, USA, September 26-28, 2011. Proceedings, Part I*, 2011, pp. 728–737.
- [27] D. Coombs, M. Herman, T. Hong, and M. Nashman, "Real-time obstacle avoidance using central flow divergence and peripheral flow," in *ICCV*, 1995, pp. 276–283.
- [28] G. Farnèbäck, "Two-frame motion estimation based on polynomial expansion," in *Image Analysis, 13th Scandinavian Conference, SCIA 2003, Halmstad, Sweden, June 29 - July 2, Proceedings*, 2003, pp. 363–370.
- [29] R. C. Nelson and Y. Aloimonos, "Obstacle avoidance using flow field divergence," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 10, pp. 1102–1106, 1989.
- [30] Y. Barniv, "Expansion-based passive ranging," NASA Ames Research Center, Moffett Field, California, Tech. Rep. 19930023159, June 1993.
- [31] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.*, vol. 17, no. 1-3, pp. 185–203, 1981. [Online]. Available: [https://doi.org/10.1016/0004-3702\(81\)90024-2](https://doi.org/10.1016/0004-3702(81)90024-2)
- [32] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proceedings of the 7th International Joint Conference on Artificial Intelligence, IJCAI '81, Vancouver, BC, Canada, August 24-28, 1981*, 1981, pp. 674–679.
- [33] P. Weinzaepfel, J. Revaud, Z. Harchaoui, and C. Schmid, "Deepflow: Large displacement optical flow with deep matching," in *IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013*, 2013, pp. 1385–1392. [Online]. Available: <https://doi.org/10.1109/ICCV.2013.175>
- [34] T. Camus, "Real-time optical flow," Ph.D. dissertation, Brown University, Department of Computer Science, 1994.
- [35] D. N. Lee, "A theory of visual control of braking based on information about time-to-collision," *Perception*, vol. 5, no. 4, pp. 437–459, 1976.
- [36] J. K. Johnson, "Selective determinism for autonomous navigation in multi-agent systems," Ph.D. dissertation, Indiana University, Department of Computer Science, 2017.
- [37] —, "A novel relationship between dynamics and complexity in multi-agent collision avoidance," in *Proceedings of Robotics: Science and Systems*, Ann Arbor, Michigan, June 2016.
- [38] J. Johnson, Y. Zhang, and K. Hauser, "Semiautonomous longitudinal collision avoidance using a probabilistic decision threshold," in *IROS 2011 International workshop on Perception and Navigation for Autonomous Vehicles in Human Environment*, 2011.
- [39] —, "Minimizing driver interference under a probabilistic safety constraint in emergency collision avoidance systems," in *2012 15th International IEEE Conference on Intelligent Transportation Systems*, Sept 2012, pp. 457–462.
- [40] A. Weinstein and M. L. Littman, "Open-loop planning in large-scale stochastic domains," in *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, ser. AAAI'13. AAAI Press, 2013, pp. 1436–1442.
- [41] C.-H. Yu, J. Chuang, B. Gerkey, G. J. Gordon, and A. Ng, "Open-loop plans in multi-robot pomdps," Stanford University, Tech. Rep., 2005.