

Predicted Composite Signed-Distance Fields for Real-Time Motion Planning in Dynamic Environments

Mark Nicholas Finean, Wolfgang Merkt, and Ioannis Havoutis

Oxford Robotics Institute, University of Oxford
 { mfinean, wolfgang, ioannis }@robots.ox.ac.uk

Abstract

We present a novel framework for motion planning in dynamic environments that accounts for the predicted trajectories of moving objects. We explore the use of *composite signed-distance fields* in motion planning and detail how they can be used to generate signed-distance fields (SDFs) in real-time to incorporate predicted obstacle motions; to achieve this, we introduce the concept of *predicted signed-distance fields*. We benchmark our approach of using composite SDFs against performing exact SDF calculations on the workspace occupancy grid. Our proposed technique generates predictions substantially faster and typically exhibits an 81–97% reduction in time for subsequent predictions. We integrate our framework with GPMP2 to demonstrate a full implementation of our approach in real-time, enabling a 7-DoF Panda manipulator to smoothly avoid a moving obstacle in simulation and hardware experiments.

1 Introduction

To integrate autonomous systems into our daily lives, we need to ensure that they operate safely in the dynamic world around them. We must, therefore, develop robots that can adapt their movements as necessary to avoid moving obstacles. This holds whether we consider robots in a household environment avoiding humans and pets, autonomous vehicles on a road avoiding other cars and pedestrians, or robots operating in a dynamic industrial setting.

There has been significant research into motion planning; however, the current ability of robots to react in dynamic environments, with moving obstacles, is still lacking. Many approaches rely on fast re-optimisation to account for changes in the environment, yet do not account for the predicted motions of the obstacles. By neglecting predicted motions, robots encounter failure cases, as depicted in Figure 2, where depending on the speed of the obstacle, the robot repeatedly plans to move in front of the moving obstacle. This scenario occurs because the optimisation can become trapped in a local minimum from which it cannot escape since it does not have full information of the obstacle trajectory during the early iterations. The resultant trajectory is both sub-optimal and potentially dangerous. In this paper,

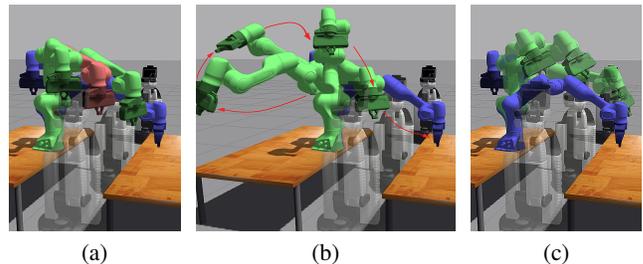


Figure 1: Real-time collision avoidance in a dynamic environment using our proposed framework. The 7-DoF Panda arm is performing a pick-up task, positioning the end-effector over a gap onto another table. During execution, a Toyota HSR travels along the gap, acting as a moving obstacle. Poses are shown in blue (start and goal configurations), green (collision-free), or red (collision). 1a) Without environment updates during execution. 1b) With environment updates and trajectory re-optimisation during execution. 1c) Our method - including obstacle trajectory prediction. We incorporate ESDF predictions which are generated at 400 Hz. The update loop runs at 3.7 Hz at the start and 10 Hz for the remainder of the trajectory.

we show that the solution to this is to incorporate predicted obstacle motions into the planning problem.

Motion planners that would otherwise use signed-distance fields are forced to adopt alternative methods, such as binary collision costs (Park, Pan, and Manocha 2012), to incorporate predictions. This is because the computation time for calculating SDFs prohibits real-time usage—i.e., the ability to generate predictions at a much higher frequency than environment observations are made, typically 30 Hz. Our results motivate using predictions and show that we can generate real-time predictions for signed-distance fields by using *composite signed-distance fields*.

We integrate our method of composite SDF predictions into our novel framework which uses the GPMP2 motion planner (Mukadam et al. 2018) to solve for collision-free trajectories in environments with moving obstacles.

The key contributions of this paper are:

1. Introduction of *predicted signed-distance fields* for mo-

- tion planning in the presence of moving obstacles.
2. Extension of GPMP2 to plan using predicted obstacle trajectories.
 3. A novel framework for motion planning in dynamic environments using composite signed-distance fields.

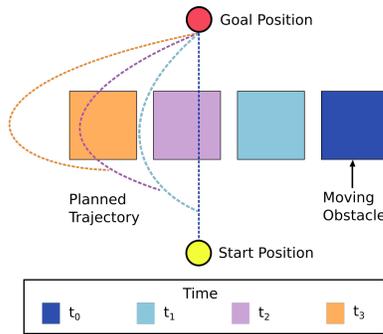


Figure 2: Toy example showing the inherent limitation of the execute-and-update method common in trajectory optimisation: The task is to plan a collision-free trajectory to reach the goal state in the presence of moving obstacles. During execution, we continue to observe the latest position of the moving obstacle and update the planned trajectory. As an obstacle approaches the intersection with the current planned trajectory, this repeated optimisation can result in the trajectory plan chasing in front of the obstacle.

2 Related Work

Motion planning in the presence of moving obstacles is an open problem, while solutions generally fall into two categories. The first is to assume full prior knowledge of the moving obstacle trajectories in the scene (Menon, Cohen, and Likhachev 2014; Hsu et al. 2002; Merkt, Ivan, and Vijayakumar 2019). The second approach is ‘continuous replanning’ in which the motion planner either re-optimises and adapts the current planned trajectory, or considers multiple trajectory modes at any one time, such as in ITOMP (Park, Pan, and Manocha 2012), and smoothly switches between them as new information is provided (Kolur et al. 2019; Zucker et al. 2013). However, there have been few works that incorporate predicted obstacle motions into the motion planning of articulated systems.

Nam *et al.* model obstacle motions as a random walk process and assign an artificial potential based on the probability density function of their predicted positions (Yun Seok Nam, Bum Hee Lee, and Moon Sang Kim 1996). Park *et al.* proposed IAMP which integrates ITOMP with human motion prediction based on supervised learning (Park, Park, and Manocha 2019). Li and Shah implement a probabilistic roadmap approach in which the roadmap is constructed using an obstacle motion prediction, essentially assuming a prior knowledge of the trajectory (Li and Shah 2019).

To further explore this problem, we first consider the environment representations used in motion planners.

Binary occupancy information enables search-based methods such as A* and D*-Lite, as well as sampling-based methods like Probabilistic Road Maps (PRMs)

and Rapidly-exploring Random Trees (RRTs), to perform collision avoidance (Hart, Nilsson, and Raphael 1968; Koenig and Likhachev 2002; Kavraki et al. 1996; Kavraki, Kolountzakis, and Latombe 1998; LaValle 1998; Kuffner and LaValle 2000). In contrast, continuous optimisation-based approaches additionally require gradients so commonly use a Euclidean distance transform (EDT), or Euclidean signed-distance field (ESDF), to represent the environment.

When planning in 3D space, the environment is most often split into an array of discrete volume elements. We define a ‘voxel’ to represent a single discrete volume in space. An ESDF is a 3D voxelgrid in which each cell contains the signed-distance to the closest obstacle surface. The sign denotes whether the cell is within or outside the surface boundary. In addition to providing distance information which can be used to assign collision costs, tri-linear interpolation enables us to obtain gradients.

Motion planners such as CHOMP, TrajOpt, and GPMP2 use ESDFs to represent the environment and demonstrate them to be very effective in enabling fast planning in static environments (Ratliff et al. 2009; Schulman et al. 2014; Mukadam et al. 2018). However, the limitations of SDFs become apparent when we look to apply them in dynamic environments.

ITOMP considered obstacle trajectory prediction for use in dynamic environments. However, in contrast to their general approach for static collision costs, which uses EDTs, they implement a simple occupancy cost for predicted motions based on “geometric collision detection between the robot and moving obstacles” (Park, Pan, and Manocha 2012). At each time-step in the optimisation a binary cost is allocated for every moving obstacle that is in collision. The reason for their approach is that EDTs are commonly considered to have CPU compute times that are not fast enough for real-time performance, thus are pre-computed and assumed to be static.

Various methods have been proposed to reduce the compute time for exact EDTs. Maurer *et al.* show that an exact distance transform can be calculated in linear time (Maurer, Qi, and Raghavan 2003). Approximate methods, such as the Chamfer Distance Transform (CDT) and Fast Marching Method (FMM), have also been proposed (Jones, Baerentzen, and Sramek 2006) to achieve fast approximations. In robotics, we often accumulate information over time; this is in contrast to areas such as medical image processing, where an ESDF will be constructed from all of the information available. This has led to significant work being conducted into faster integration of new occupancy information into signed-distance fields. Oleynikova *et al.* found that Truncated Signed-Distance Fields (TSDFs) are faster to construct than Octomaps, a widely used package that maps the environment using octrees (Hornung et al. 2013), and proposed, VoxBlox, a method to build an ESDF from projective TSDFs incrementally (Oleynikova et al. 2017). In contrast to ESDFs, TSDFs use a ‘projected distance’ metric and consider distances only within a truncated radius of surface boundaries. TSDFs are commonly used in vision for surface reconstruction (Newcombe et al. 2011; Whelan, Kaess, and

Fallon 2012).

Despite the speed-ups achieved by state-of-the-art methods, they are not able to calculate SDFs fast enough for use in a real-time motion planner that uses what we term ‘*predicted signed-distance fields*’ - predictions of what the ESDF of the workspace will be for future times based on estimates of obstacle geometry, velocity, and direction. We propose using *composite signed-distance fields* to generate predicted signed-distance fields in real-time.

Zucker *et al.* first mentioned the concept of composite SDFs, noting that distance fields are compositional under the \min operation (Zucker et al. 2013). Thus the distance field computation can be reduced to a minimisation across a set of pre-computed distance field primitives. In the computer graphics community, composition of object-centric SDFs is used to unionise objects in ‘Constructive Solid Geometry’ as well as to combine the SDFs of individual objects into a single SDF of the scene for ray-marching (Reiner, Mückl, and Dachsbacher 2011).

Zucker *et al.* use composite SDFs in ‘CHOMP-R’ to quickly update their environment representation in response to a changing environment. However, to our knowledge, they have not been used in a predictive manner. We provide further analysis of this technique and leverage it in our novel framework which extends it to the domain of integrating predicted object trajectories into motion planning.

To demonstrate our framework, we integrate it with GPMP2 which formulates motion planning as probabilistic inference on a factor graph and uses Gaussian Processes (GPs) to represent continuous-time trajectories with a small number of states (Mukadam et al. 2018). Trajectories are found by using the GTSAM framework (Dellaert 2012) to exploit sparsity in the problem and perform fast probabilistic inference on factor graphs. Collision avoidance is applied via the use of ‘obstacle factors’, which evaluate the collision-free likelihood, using a hinge-loss obstacle cost for each time-step. The hinge-loss is calculated, in both CHOMP and GPMP2, by using an approximate robot model, comprised of spheres, to query the signed-distance field (Ratliff et al. 2009; Mukadam et al. 2018). However, Mukadam *et al.* use a single pre-computed SDF of the surrounding environment that is shared by all of the obstacle factors (Mukadam et al. 2018) and so does not incorporate the dynamics of moving obstacles. Kolar et al. consider dynamic environments but update all of the obstacle factors to use the same updated SDF and then re-optimize to solve for an updated trajectory plan (Kolar et al. 2019).

3 Methods

Part 3.1 describes our method of predicting trajectories for dynamic obstacles. Part 3.2 details our novel framework for predicting environment representations, using composite signed-distance fields, and then outlines how they can be integrated with a motion planner. We describe an experiment to examine the speed-up attained by using composite signed-distance fields in contrast to exact computations. In Part 3.3, we describe how we adapted the GPMP2 motion planner to react in dynamic environments as well as an experiment which demonstrates its applicability. In Parts 3.4

and 3.5, we demonstrate our fully integrated framework in simulation and on hardware, respectively. In both, a 7-DoF Panda manipulator successfully avoids a moving obstacle.

3.1 Obstacle Trajectory Prediction

To demonstrate that our approach enables the incorporation of obstacle trajectory predictions, we consider a simple prediction strategy for the occupancy of the workspace. Consider a 3D workspace, \mathcal{W} , which comprises of a set of $i = 1, \dots, n$ stationary objects, \mathcal{O}_i^s , and $j = 1, \dots, m$ moving (dynamic) objects, \mathcal{O}_j^d . At a given time, t , a moving object with a centroid position, $\mathbf{x}_j(t)$, will travel at velocity $\mathbf{v}_j(t)$. We write the collection of discretised positions and velocities as $\mathbf{X}(t) = [\mathbf{x}_1(t), \mathbf{x}_2(t), \dots, \mathbf{x}_m(t)]^T$ and $\mathbf{V}(t) = [\mathbf{v}_1(t), \mathbf{v}_2(t), \dots, \mathbf{v}_m(t)]^T$.

In our simulated experiments, we assume noiseless sensor data and subsequently an accurate occupancy grid of the workspace, $\mathcal{G}(t)$, at each time-step. At any given time, we can identify and isolate occupied regions of the workspace. From these regions, we calculate the centroid and list of voxels associated with each object. By comparing two successive frames, at times t and $t + \delta t$, we identify which objects are moving and estimate their respective velocities from the motion of each centroid—this enables us to identify occupancy grids associated with the static scene, \mathcal{G}_{static} , and the moving obstacles, $\mathcal{G}_{moving}(t)$.

With centroid and velocity information, we propagate each object using a constant-velocity model to produce a predicted trajectory of the centroid. This method inherently gives accurate trajectory predictions when an object is moving in a straight line with constant-velocity. However, we emphasise that the prediction model is fully replaceable and that a more sophisticated object tracking and trajectory prediction can be implemented within our framework, e.g. a KLT tracker or Unscented Kalman Filter (Jianbo Shi and Tomasi 1994; Raj et al. 2016).

3.2 Composite Signed-Distance Field Generation

Building on the previous description of how to identify the static and moving parts of the scene, we use \mathcal{G}_{static} to calculate the static SDF \mathcal{S}_{static} . Next, we consider an occupancy box, \mathcal{B}_j , around each of the m moving objects, \mathcal{O}_j^d , that we have identified in \mathcal{G}_{moving} . Each occupancy box is chosen to be of equal size to \mathcal{O}_j^d , plus an additional ϵ in all directions. For each \mathcal{B}_j , we calculate the corresponding object SDF, \mathcal{S}_j . The resultant SDF is the smallest cuboid which encloses all voxels in 3D space with signed-distance values equal to or less than a distance ϵ away from the surface of the object. We introduce ϵ as a safety margin desired to be penalised in collision avoidance. When this method is integrated with a motion planner, ϵ should equal the desired safety margin around objects to be penalised plus the size of the largest collision sphere associated with the approximate robot model.

Similar to $\mathbf{X}(t)$ and $\mathbf{V}(t)$, we denote the stacked vector of moving SDFs at a given time as $\mathcal{S}(t)$. In practice, we may accumulate more information about the shape of each moving object, hence $\mathcal{S}(t)$ is a function of time. However, for demonstration, our simulations assume noiseless sensor

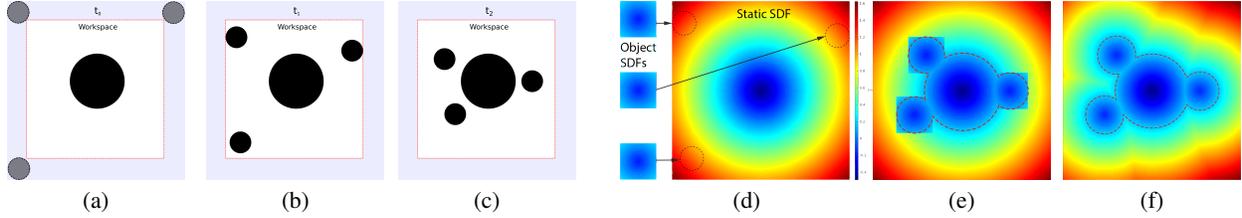


Figure 3: 3a, 3b and 3c depict a toy example of three moving spheres entering the workspace in which a static obstacle (large central sphere) is present. White regions in the workspace represent free space, while black regions are occupied. 3d shows an SDF of the static environment, with the tracked objects in 3b overlaid for illustration purposes. For each of the tracked objects in the scene, an SDF is calculated and associated with them. By tracking the positions of the moving objects, we infer their velocities in order to make predictions of their future positions. At these future positions, we can superpose the object SDFs onto the static SDF using a \min operation—the result is a composite SDF. 3e shows a composite SDF for t_2 . For comparison, 3f shows the corresponding exact SDF. Critically for motion planning, the two are identical for distances up to ϵ away from the obstacle surface boundaries, as indicated by the red dashed line.

Algorithm 1 Real-time Update Loop

Usage: Continuously updates the motion planner

Initialisation :

1: $\mathcal{S} = \text{getObjectSDFs}()$

2: $\mathcal{S}_{static} = \text{getStaticSDF}()$

Execution :

3: **while** not shutdown **do**

4: $\mathcal{S} = \text{updateObjectSDFs}()$

5: $\mathbf{X} = \text{updateObjectPositions}()$

6: $\mathbf{V} = \text{updateObjectVelocities}()$

7: **for** t in *times to predict* **do**

8: $\mathbf{X}_{predicted} = \text{calcFutureObjPositions}(t)$

9: $\mathcal{S}_{predicted} = \text{predictSDF}(t, \mathbf{X}, \mathbf{V}, \mathcal{S}, \mathcal{S}_{static})$

10: $\text{updateMotionPlanner}(t, \mathcal{S}_{predicted})$

11: **end for**

12: $\text{reoptimiseAndSendTrajectory}()$

13: **end while**

data and so do not acquire additional shape information in subsequent steps; we thus refer to $\mathcal{S}(t)$ simply as \mathcal{S} .

For any future time, t_f , we can predict the occupancy of the workspace by propagating the current position $\mathbf{x}_j(t_n)$ of each moving obstacle SDF \mathcal{S}_j to its predicted position using velocity $\mathbf{v}_j(t_n)$. Similar to considering an occupancy grid as the superposition of all occupied regions, we consider the workspace SDF as the superposition of all object SDFs in the scene—this forms the basis of a composite SDF.

To calculate the resultant composite SDF for a given time, we take the minimum voxel value (i.e., signed-distance) of all overlapping scene and object SDFs. By construction, the composite SDF is guaranteed to be accurate up to the distance ϵ . We illustrate this process in Figure 3.

Algorithms 1 and 2 describe how we use this method in an update loop for motion planning with SDF predictions generated in real-time. While obstacle prediction in motion planning is not novel, to our knowledge this has not been performed by predicting the SDF of the environment for different time-steps to solve for a trajectory. Our approach thus

Algorithm 2 $\text{predictSDF}(t, \mathbf{X}, \mathbf{V}, \mathcal{S}, \mathcal{S}_{static})$

Output: $\mathcal{S}_{predicted}$

1: $\mathcal{S}_{predicted} = \mathcal{S}_{static}$

2: **for** $j = 1 : m$ **do**

3: $inds = \text{getOccupiedIndices}(t, \mathbf{x}_j, \mathbf{v}_j)$

4: $\mathcal{S}_{predicted}(inds) = \min(\mathcal{S}_{predicted}, \mathcal{S}_j)$

5: **end for**

6: **return** $\mathcal{S}_{predicted}$

relates to time-configuration space planning, explicitly encoding the state of the environment at different points in time (Merkt, Ivan, and Vijayakumar 2019).

To evaluate the speed-up achieved using composite SDFs, we constructed a set of four dynamic datasets modelled on the ‘table and cabinet’ setup presented in (Mukadam, Yan, and Boots 2016). Each of the environments contains a table and cabinet with the addition of: a) one small moving box, b) two small moving boxes, c) one moving pillar, d) two moving pillars. Figure 4 shows aerial and three-quarter perspective views of the ‘single moving pillar’ setup. We also include a simpler environment consisting of an empty workspace with a large moving block.

For each dynamic environment, we simulate 31 different time-steps in a $96 \times 96 \times 96$ workspace with 4 cm resolution. At each time-step, we produce composite SDF predictions for the remaining time-steps. We perform the same simulation and benchmark against the exact SDF computations as calculated in GPMP2, using MATLAB’s `bwdist` function, on the predicted occupancy grids.

3.3 Dynamic GPMP2

We adapt the original GPMP2 implementation to enable rapid updates of the SDF associated with each obstacle factor *independently*. The process of quickly updating obstacle factors facilitates a motion planner that can re-optimize and adapt planned trajectories to a changing environment.

To examine the effect of using independent SDFs for each obstacle factor, and to assess the importance of incorporat-

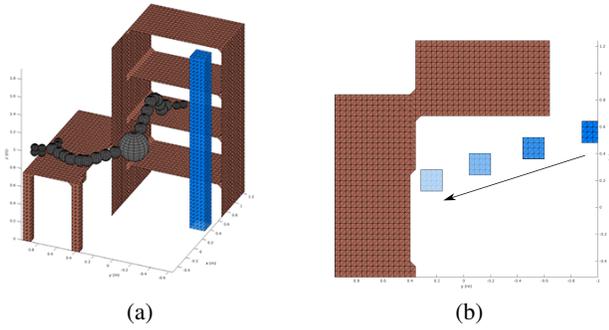


Figure 4: ‘Single moving pillar’ environment - Aerial and three-quarter perspective views of an example environment in which a tall pillar traverses the floor. Example start and goal configurations are depicted using a sphere approximation of the WAM robot.

ing obstacle dynamics, we considered and compared three different scenarios in dynamic environments. Each scenario considers a different degree of perception; they are as follows:

1. *Static*: Assumes a static workspace and all obstacle factors remain unchanged throughout the execution of the trajectory. The factor graph is thus optimised only once.
2. *Execute-and-update*: Obstacles factors are initialised with the initial SDF of the scene. During execution, at each time-step, all obstacle factors with a time-index greater than or equal to the current time are updated to use the current observation of the workspace. The trajectory is re-optimised and executed at each time-step in an iterative process.
3. *Full prior knowledge*: Each obstacle factor is provided with the known SDF of the environment at the associated time. The factor graph has maximum knowledge of the environment evolution and is thus only optimised once.

The factor graphs associated with each of the perception scenarios are illustrated in Figure 5.

To test each scenario, we used a 7-DoF WAM robot arm in simulation to plan trajectories in the ‘single moving pillar’ setup shown in Figure 4. Plans were conducted over seven different combinations of start and goal configuration, and a range of moving obstacle speeds, 0.1 m/s to 1.0 m/s. We recorded the resultant smoothness (GP) cost associated with each trajectory and the number of collisions that the robot made with the environment.

To formalise the task, we constructed a factor graph comprising of 31 time-indexed variable nodes, separated by 0.1 s intervals, to form a 3 s time horizon. Prior factors were added to specify the start and goal configurations, and Gaussian priors used to connect variable nodes, as in (Mukadam et al. 2018). We assigned an obstacle factor to each variable node and interpolated obstacle factors between pairs of variable nodes with a discretisation of $\tau = 5\text{ms}$ ($n_{int} = 19$) to provide dense collision checking along the trajectory. Further parameters used were $\epsilon = 0.2\text{ m}$ to encourage safe clearance of obstacles, as well as $\sigma_{cost} = 0.2$, as a result of tuning, to weight the relative costs of trajectory smoothness and col-

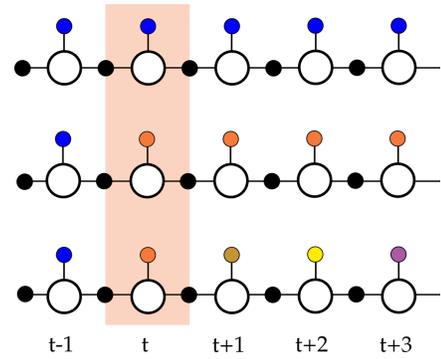


Figure 5: Factor graphs associated with three different feedback scenarios. Variable nodes are shown in white, prior factors in black, and coloured obstacle factors. *Top* - the obstacle factors share the same SDF and remain unchanged during execution. *Middle* - as the trajectory is executed, the current and all future obstacle factors are replaced with the latest SDF. *Bottom* - as the trajectory is executed, the current obstacle factor is updated to the latest observed SDF while all future obstacle factors receive updated SDF predictions.

lision avoidance. The environment was discretised at 4 cm resolution in a $96 \times 96 \times 96$ voxel workspace.

We performed another similar experiment on the experimental setup shown in Figure 1, in which a Franka Emika Panda 7-DoF robotic arm was tasked with reaching across the gap between two tables to achieve a ‘pick-up’ goal configuration. During execution, a Toyota HSR traverses the walkway between the tables, acting as a moving obstacle. For this setup, we used 42 different combinations of start and goal configuration. The moving obstacle speed was chosen to be in the range of 1.0 m/s to 2.0 m/s, most closely resembling walking speeds of a human. We chose $\tau = 0.05\text{ s}$ ($n_{int} = 1$) to save compute time and because the time interval between variable nodes was already small. Other parameters used were $\sigma_{cost} = 0.05$ and $\epsilon = 0.3\text{ m}$ based on cost tuning. As before, the environment was discretised at 4 cm resolution in a $96 \times 96 \times 96$ voxel workspace.

3.4 Closed-loop Implementation

We combined our predictive framework using composite signed-distance fields with our dynamic approach, applied within GPMP2, to solve planning problems in dynamic environments. The control loop is shown in Algorithm 1. In each iteration of the control loop, we observe the environment and update the workspace occupancy grid. We populate the occupancy grid by tracking the position of model states in the Gazebo simulation and inserting occupied regions into the occupancy grid at corresponding locations.

Using the methods discussed in Section 3.1, we segment and track obstacles online. A constant-velocity model is applied to each of the segmented moving objects, enabling us to predict the future locations of each object and generate composite SDFs for all obstacle factors with time-index equal to or greater than the current time. To our knowledge, this is the first work to implement and update different SDFs

for each time-step as in time-configuration space planning. A prior factor is also added for the current time-index to set the current configuration of the robot. The factor graph is then re-optimised and the new trajectory is merged with the current trajectory being executed on the robot.

We demonstrate our closed-loop implementation of the proposed framework on a Panda 7-DoF robotic arm in Gazebo, using the same task as described previously in which the Panda reaches across a gap to place the end-effector over a table, as in a pick-and-place task. We set the HSR base movement speed at 1.4 m/s to test against a speed comparable with the average walking pace (Hermodsson et al. 1994; Nolan et al. 2018; Wang et al. 2019).

3.5 Hardware Experiments

We validate our dynamic GPMP2 approach in the presence of uncertainty in the execution of the trajectory, as well as the pose of the dynamic obstacle, in hardware experiments. We use a Franka Emika Panda arm in conjunction with a Realsense D435i RGB-D camera for object tracking and detection. We used *apriltag_ros* (Malyuta et al. 2019) to identify the pose of an AprilTag (Wang and Olson 2016) fiducial marker affixed to a dynamic obstacle (box). For safety purposes, we used a longer time-horizon of 5 s and an epsilon value of 0.4 m to encourage both lower joint speeds and comfortable collision avoidance. The robot was positioned on a single table and tasked with a variety of motions, including the same motion as in the previous experiment, and avoiding the box as moved by the operator.

4 Evaluation

We conducted experiments using an 8-core Intel Core i7-9700 CPU @ 4.50 GHz and 2133 MHz DDR4 RAM. Signed-distance field computations were performed in series, as in the original GPMP2 implementation, using the `bwdist` function in MATLAB. GTSAM was configured to use multi-threading enabled for linearising the factor graph.

4.1 Composite SDF Prediction

Computation times for composite signed-distance fields and full computation from $96 \times 96 \times 96$ occupancy grids are shown in Table 1.

The composite initialisation time, highlighted in grey, comprises of the time to calculate the SDFs associated with the static workspace, \mathcal{S}_{static} , and each of the moving obstacles in the scene, \mathcal{S} . This initial time penalty must be paid at the startup of the experiment or any subsequent time you wish to update the object and static environment SDFs. However, our results show that beyond this initialisation penalty, it is significantly faster to generate predicted composite SDFs than to perform full SDF calculations; we see a 7.9 x to 28.0 x speed-up for subsequent predictions depending on the discretisation of the workspace. This has a significant impact on enabling SDF predictions to be generated in real-time—conducting our experiment for a $96 \times 96 \times 96$ discretisation of the workspace, as in our closed-loop implementation, gives a mean SDF generation rate of 540 Hz, compared with 36 Hz for the benchmark.

4.2 Dynamic GPMP2

The performance of the three scenarios described in Section 3.3, on the ‘single moving pillar’ dataset, are presented in Figure 6. The figure shows an absence of data for the obstacle speed of 0.5 m/s; at 0.5 m/s, the obstacle intersects the goal configurations at the end of the 3 s time horizon, resulting in inevitable collisions. Therefore, we exclude results for 0.5 m/s from our analysis. Of 69 planning results, we exclude three for which none of the planning scenarios found a collision-free trajectory.

Our results strongly support the incorporation of future obstacle motions into motion planning. Figure 6 shows that in most cases, including the obstacle trajectory motion into the planning significantly increases the likelihood of both finding a collision-free trajectory and reducing the jerk associated with the trajectory. Furthermore, the results motivate the use of our proposed framework in which we generate predicted ESDFs in real-time.

We note here that our implementation of a dynamic obstacle factor in GPMP2 incurs a cost of ≈ 2 ms to be replaced by an obstacle factor with a different SDF.

4.3 Closed-Loop Implementation

Our closed-loop implementation was successful in providing re-optimised trajectories fast enough for a Panda arm to avoid a moving HSR robot while reaching across a gap between two tables. Figure 1 illustrates the trajectory taken in each of the three cases. The ‘static’ case would result in a collision and the ‘execute and update’ would exhibit erratic movements. Figure 7 shows a comparison of the resultant joint position trajectories for the ‘execute and update’ against our composite prediction method. The results lend further support that by including *predicted SDFs* into motion planning, we can obtain smoother robot trajectories.

This experiment was conducted using a $96 \times 96 \times 96$ workspace with 4 cm resolution. The parameters for the motion planner were $n_{int} = 1$, $\sigma_{obs} = 0.05$, $\epsilon = 0.3$ m. A time horizon of 3 s was given, with variable factors time-indexed in steps of $\delta_t = 0.1$ s, resulting in 31 support states. The first iteration of the control loop results in the slowest update frequency because more future factors need to be updated and more predictions are generated. Averaged over 100 runs, our implementation runs at 3.4 Hz in the early iterations, with a standard deviation of 0.1 Hz. In the later iterations, we achieve an update frequency of ≈ 10 Hz. We demonstrate the resultant trajectories for each of the three cases in the accompanying video.¹

4.4 Hardware Experiments

Our implementation performed well on the real robot arm, producing re-planned trajectories which successfully avoided the dynamic object in many cases. An example comparison against using composite SDFs but without prediction is shown in Figure 8. By accounting for the predicted motion of the obstacle, the manipulator is able to avoid the obstacle by following a much smoother trajectory.

¹<https://youtu.be/vJH8qBRMedw>

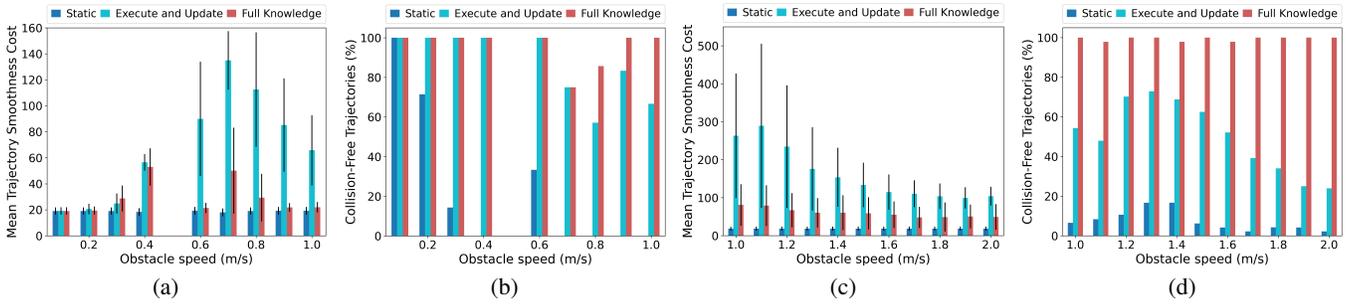


Figure 6: Results for experiments conducted on the WAM arm in the ‘single moving pillar’ setup are shown in Figures 6a and 6b. Results for the Panda arm avoiding a moving HSR are shown in 6c and 6d. We see that in all cases, using full knowledge of obstacle trajectories provides at least an equal chance of finding a collision-free trajectory using GPMP2; for most obstacle speeds, the improvement is significant.

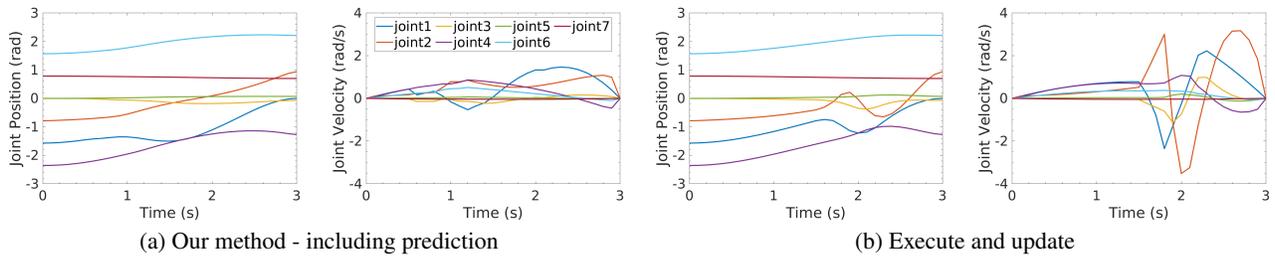


Figure 7: Example trajectories for joint position and velocity. Our results highlight that the inclusion of environment prediction in motion planning can lead to smoother trajectories which exhibit less erratic movements.

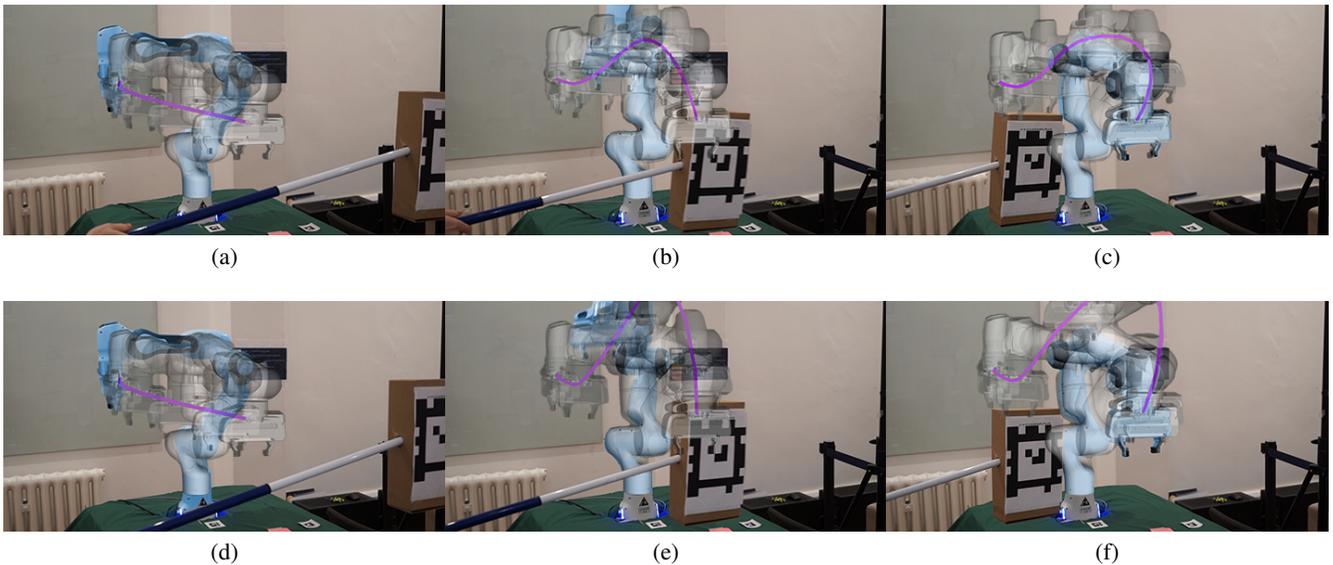


Figure 8: *Top row* - online motion planning using our approach of composite signed distance fields and prediction. *Bottom row* - using composite signed distance fields but without prediction. From left to right, the images show the planned trajectory for times $t = 0$ s, $t = 2.5$ s and $t = 5.0$ s (resultant trajectory). The blue highlight indicates the configuration of the robot at the current time-step while the purple line show the planned end-effector trajectory. The predictive component of our framework enables a smoother resultant trajectory, as shown in Figure 8c.

	Workspace Size (Voxels Per Side)								
	64	96	128	160	192	224	256	288	320
Composite Init. (ms)	13 ± 1	31 ± 4	65 ± 3	126 ± 6	211 ± 3	329 ± 6	475 ± 7	659 ± 5	936 ± 23
Full Computation (ms)	9.3 ± 1.0	28 ± 2	62 ± 3	121 ± 5	217 ± 11	338 ± 9	493 ± 13	687 ± 18	960 ± 28
Composite Prediction (ms)	0.6 ± 0.6	1.9 ± 0.2	4.0 ± 0.9	7.4 ± 0.5	24 ± 6	42 ± 2	62 ± 3	86 ± 3	121 ± 3
Repeat Prediction Speed-Up	15.5x	28.0x	15.5x	16.4x	9.0x	8.0x	8.0x	8.0x	7.9x

Table 1: Signed-Distance Field Timing Analysis

We found that while neglecting to predict the obstacle motion was also effective at avoiding obstacles due to the high re-plan frequency, the robot would exhibit erratic behaviours, in particular when the obstacle position coincided with the desired arm goal state.

5 Discussion

Our proposed framework gives compelling results in support of incorporating predicted obstacle trajectories into motion planning in time-configuration space. While we implemented a constant-velocity model to demonstrate our framework’s novel ability to include obstacle trajectory predictions directly into SDF-based motion planning, one can replace this with a more complex model that takes into account any prior information we have on the obstacle’s trajectory, including its history. Exploring different prediction models for obstacle motion was beyond the scope of this work and as such we did not explore the avoidance of obstacles that follow complex motions, such as curved trajectories; we leave this for future work. Without using a more complex model, our current implementation is likely to fail in circumstances where the obstacle suddenly changes direction or stops.

A limitation of the proposed method is the potentially large memory requirements associated with assigning and storing different signed-distance fields for different time-indexed obstacle factors. For example, if a $300 \times 300 \times 300$ discretisation of the workspace is required for an application, our previous problem of a 3 s time horizon, discretised into 31 time-steps, will require ≈ 6.24 GB RAM.² We thus encounter a trade-off; minimising the workspace size results in quicker computation and lower memory requirements, however, to monitor a larger workspace requires the resolution to be decreased. We believe that this method could be adapted in future work to perform the superposition on a query basis, rather than computing the composite signed-distance field for the entire workspace.

Generating composite predictions for each time-step, as well as the updating of the obstacle factors, are operations which seem suitable for parallelisation. By updating factors in parallel, we believe that significantly higher update rates could be achieved allowing for more interactive and responsive behaviour.

We have demonstrated benefits of predicted composite SDFs that are independent of state-of-the-art SDF computing techniques. In our future work, we will evaluate dif-

²A double is stored using 8 bytes. $300^3 \times 8 = 2.16 \times 10^8$ bytes = 206 MB. $31 \times 206 \approx 6.24$ GB

ferent methods that can leverage our technique and framework, particularly parallel implementations. We also note that while pre-computing a static scene is acceptable for a fixed-base robot arm, it is less useful in the context of a mobile robot performing tasks in an unmapped region of space. To address this, we are working on integrating our approach with ESDFs that are updated from live point-cloud data.

6 Conclusion

This paper explored the application of composite signed-distance fields to motion planning in dynamic environments. We first exploited the speed of the `min` operation for composition of SDFs, and secondly, that signed-distance field representations must only be accurate up to a specified ϵ for motion planning problems. We show that composite SDFs can be used to provide significant speed-up for generating SDFs when accounting for moving obstacles. We investigated motion planning with a GPMP2 implementation which uses dynamic obstacles factors, enabling the planner to account for moving obstacles. Over a range of tasks and two different robot platforms, our results show that by incorporating predicted obstacle trajectories, we can significantly reduce the smoothness cost of trajectories and the rate of collisions. We leveraged composite SDFs and a dynamic GPMP2 implementation to present a novel framework which exploits the sparsity of the workspace and the compositional nature of signed-distance fields to generate real-time predictions of the workspace SDF, *predicted signed-distance fields*. We verified our approach on a real 7-DoF Panda arm with obstacle tracking performed using fiducial sensing, demonstrating that it can plan trajectories in dynamic environments and successfully avoid moving obstacles.

Acknowledgments

This research was supported by (1) the UK Engineering and Physical Sciences Research Council (EPSRC) through the University of Oxford Centre for Doctoral Training, Autonomous Intelligent Machines and Systems (AIMS, grant references EP/L015897/1 and EP/R512333/1), the UK RAI Hub for Offshore Robotics for Certification of Assets (ORCA, grant reference EP/R026173/1), and the EPSRC grant Robust Legged Locomotion (EP/S002383/1), as well as (2) the European Commission under the Horizon 2020 project Memory of Motion (MEMMO, project ID: 780684).

References

Dellaert, F. 2012. Factor Graphs and GTSAM: A Hands-on Introduction. Technical report, Georgia Tech.

- Hart, P.; Nilsson, N.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Trans. on Systems Science and Cybernetics* 4(2): 100–107.
- Hermodsson, Y.; Ekdahl, C.; Persson, B. M.; and Roxendal, G. 1994. Gait in male trans-tibial amputees: A comparative study with healthy subjects in relation to walking speed. *Prosthetics and Orthotics International* 18(2): 68–77.
- Hornung, A.; Wurm, K. M.; Bennewitz, M.; Stachniss, C.; and Burgard, W. 2013. OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees. *Autonomous Robots* 34: 189–206.
- Hsu, D.; Kindel, R.; Latombe, J. C.; and Rock, S. 2002. Randomized kinodynamic motion planning with moving obstacles. *The Int. J. of Rob. Research* 21(3): 233–255.
- Jianbo Shi; and Tomasi. 1994. Good features to track. In *Proc. IEEE CVPR*, 593–600. IEEE Comput. Soc. Press.
- Jones, M.; Baerentzen, J.; and Sramek, M. 2006. 3D distance fields: a survey of techniques and applications. *IEEE Trans. on Visualization and Computer Graphics* 12(4): 581–599.
- Kavraki, L.; Kolountzakis, M.; and Latombe, J.-C. 1998. Analysis of probabilistic roadmaps for path planning. *IEEE Trans. on Rob. and Autom.* 14(1): 166–171.
- Kavraki, L.; Svestka, P.; Latombe, J.-C.; and Overmars, M. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. on Rob. and Autom.* 12(4): 566–580.
- Koenig, S.; and Likhachev, M. 2002. Improved fast replanning for robot navigation in unknown terrain. In *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, 968–975.
- Kolur, K.; Chintalapudi, S.; Boots, B.; and Mukadam, M. 2019. Online Motion Planning Over Multiple Homotopy Classes with Gaussian Process Inference. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2358–2364.
- Kuffner, J.; and LaValle, S. 2000. RRT-connect: An efficient approach to single-query path planning. In *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, 995–1001.
- LaValle, S. M. 1998. Rapidly-exploring random trees : a new tool for path planning. Technical report, Iowa State University.
- Li, S.; and Shah, J. A. 2019. Safe and efficient high dimensional motion planning in space-time with time parameterized prediction. In *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, 5012–5018.
- Malyuta, D.; Brommer, C.; Hentzen, D.; Stastny, T.; Siegwart, R.; and Brockers, R. 2019. Long-duration fully autonomous operation of rotorcraft unmanned aerial systems for remote-sensing data acquisition. *Journal of Field Robotics* .
- Maurer, C. R.; Qi, R.; and Raghavan, V. 2003. A linear time algorithm for computing exact Euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 25(2): 265–270.
- Menon, A.; Cohen, B.; and Likhachev, M. 2014. Motion planning for smooth pickup of moving objects. In *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, 453–460.
- Merkt, W.; Ivan, V.; and Vijayakumar, S. 2019. Continuous-Time Collision Avoidance for Trajectory Optimization in Dynamic Environments. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 7248–7255.
- Mukadam, M.; Dong, J.; Yan, X.; Dellaert, F.; and Boots, B. 2018. Continuous-time Gaussian process motion planning via probabilistic inference. *The Int. J. of Rob. Research* 37(11): 1319–1340.
- Mukadam, M.; Yan, X.; and Boots, B. 2016. Gaussian Process Motion planning. In *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, 9–15. IEEE.
- Newcombe, R. A.; Fitzgibbon, A.; Izadi, S.; Hilliges, O.; Molyneaux, D.; Kim, D.; Davison, A. J.; Kohi, P.; Shotton, J.; and Hodges, S. 2011. KinectFusion: Real-time dense surface mapping and tracking. In *Proc. IEEE ACM Int. Symp. Mixed Augment. Reality (ISMAR)*, 127–136.
- Nolan, C. M.; Kon, S. S. C.; Patel, S.; Jones, S. E.; Barker, R. E.; Polkey, M. I.; Maddocks, M.; and Man, W. D.-C. 2018. Gait speed and pedestrian crossings in COPD. *Thorax* 73(2): 191–192.
- Oleynikova, H.; Taylor, Z.; Fehr, M.; Siegwart, R.; and Nieto, J. 2017. Voxblox: Incremental 3D Euclidean Signed Distance Fields for on-board MAV planning. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 1366–1373.
- Park, C.; Pan, J.; and Manocha, D. 2012. ITOMP: Incremental trajectory optimization for real-time replanning in dynamic environments. In *Proc. Int. Conf. Automat. Planning and Scheduling (ICAPS)*, 207–215.
- Park, J. S.; Park, C.; and Manocha, D. 2019. I-Planner: Intention-aware motion planning using learning-based human motion prediction. *The Int. J. of Rob. Research* 38(1): 23–39.
- Raj, A.; Sivaraman, A.; Bhowmick, C.; and Verma, N. K. 2016. Object tracking with movement prediction algorithms. In *Proc. IEEE Int. Conf. Ind. Inf. Syst. (ICIIS)*, 285–290.
- Ratliff, N.; Zucker, M.; Bagnell, J. A.; and Srinivasa, S. 2009. CHOMP: Gradient optimization techniques for efficient motion planning. In *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, 489–494.
- Reiner, T.; Mückl, G.; and Dachsbacher, C. 2011. Interactive modeling of implicit surfaces using a direct visualization approach with signed distance functions. Technical Report 3, Karlsruhe Inst. of Tech.
- Schulman, J.; Duan, Y.; Ho, J.; Lee, A.; Awwal, I.; Bradlow, H.; Pan, J.; Patil, S.; Goldberg, K.; and Abbeel, P. 2014. Motion planning with sequential convex optimization and convex collision checking. *The Int. J. of Rob. Research* 33(9): 1251–1270.
- Wang, J.; and Olson, E. 2016. AprilTag 2: Efficient and robust fiducial detection. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 4193–4198.
- Wang, Z.; Liu, R.; Wu, X.; and Liu, Z. 2019. *Estimation Models for the Safety Level of Indoor Space Pedestrian Flows*, 9–18. Springer.
- Whelan, T.; Kaess, M.; and Fallon, M. 2012. Kintinuous: Spatially extended kinectfusion. *RSS Workshop on RGB-D: Advanced Reasoning with Depth Cameras* 7.
- Yun Seok Nam; Bum Hee Lee; and Moon Sang Kim. 1996. View-time based moving obstacle avoidance using stochastic prediction of obstacle motion. In *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, volume 2, 1081–1086.
- Zucker, M.; Ratliff, N.; Dragan, A. D.; Pivtoraiko, M.; Klingensmith, M.; Dellin, C. M.; Bagnell, J. A.; and Srinivasa, S. S. 2013. CHOMP: Covariant Hamiltonian optimization for motion planning. *The Int. J. of Rob. Research* 32(9-10): 1164–1193.