Planning for Robot Recovery with Varying Model Fidelity

Breelyn Kane Styler
Robotics Institute, Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
breelynk@cs.cmu.edu

For robots to be usable in the future, they need to operate dependably long term. Failure reduction for robot platforms is very important for robot usability. Robots that repeat the same failed action over and over, degrade our faith in their intelligence. If robots are going to be useful working with humans, they need to be trustworthy. For example, imagine a person having to stay at home to care for a sick relative. The robot provides more independence to the caregiver by performing routine chores that allow the caregiver to focus time on meaningful patient tasks. The robot would be expected to perform their assistive tasks without needing constant failure monitoring. In this case, the robot's reliable autonomy directly impacts the caregiver's autonomy.

This thesis investigates how a robot makes decisions to recover from failures for reliable and autonomous operation. A common approach is to enumerate hard coded recovery strategies for failures that occur frequently. However, trying to determine what failures could occur during planning is a difficult problem when there are rare failures that have not been anticipated. An alternate approach is to try and learn recovery strategies for such failures, but it is hard to imagine how one could have a robot figure out how to deal with completely unmodeled situations. We take a middle approach. We assume that the robot has the necessary knowledge, but that making use of that knowledge all the time is not feasible. The entire planning space represents the robot's knowledge base. This space contains multiple models of varying fidelity that, when combined, approximate the continuous highest possible real-world model. In this paradigm, a robot's ability to make decisions is only as good as its models. Therefore, the robot's goal is to decide at any point during its execution what the most appropriate level of fidelity is for the given task.

We assume these models are given. The planning models describe state and action information explicit to the robot's interaction with the environment. Leveraging these models, the robot is able to generate plans in different spaces that circumvent the failures it encounters. Our recovery approach focuses on what model to select and how to localize planning around a failure. For example, a robot may begin planning in a flat two dimensional space that is common for its current task. A failure occurs when the robot hits an overhang in three-dimensional space that it was not initially considering. Considering a model with the overhang height could have prevented that failure.

If the robot planned in the highest fidelity space possible, its failure rate would be reduced, so why not always use the highest applicable model? Unfortunately, necessary information may not be available for reasoning in a higher model. Limitations in the robot's sensor accuracy, and very dynamic environments both increase uncertainty. This uncertainty represents a lack of knowledge for the robot, which may not become available to the robot until later during plan execution. Additionally, failures tend to provide the robot with information. This is why it is important for the robot to start planning in a space where it is certain about information. In some cases this might be a model that explicitly reasons about uncertainty, and as unexpected events occur, the robot can switch to models with reduced uncertainty. Secondly, higher fidelity models often require more information (such as the coefficient of friction of the terrain) that is not known. Without knowing this information, the value in planning in the higher fidelity model is lost. Either way, switching resolves ambiguity by focusing model complexity to areas where the robot has the most certainty.

Switching also has the added effect of combining the robustness of a high-fidelity plan with the efficiency of an abstract representation by using higher fidelity models only when necessary. Planning in the highest fidelity space may not be necessary at all points during the robots task. For example, if the robot is moving through large unconstrained open spaces that do not require many turns, a lower fidelity model that only considers geometric planar motion might suffice. For that reason, the robot's non-uniform execution space suggests plan savings by leveraging multiple models. Our work attempts to stay as long as possible in the lowest fidelity model applicable in order to decrease plan computation time without sacrificing execution quality.

The thesis will explore topics for making smarter decisions on what models to switch to. This will happen by focusing on three areas. The first is having a principled way to organize the models in a hierarchy. Model organization is important for knowing how the models relate to each other. These relations contribute to more informed decisions such as which subsets of the model graph may be more important than others for the given task and failure. The second will address how uncertainty fits into the different models, and reasoning about model choices that involve uncertainty in representations of states or action transitions. Finally, the last part of the thesis will focus on how to make more probabilistic switching decisions. Uncertainty can generate plans that are more conservative which introduces the notion of risk with model selection.

Building a principled hierarchical graph involves breaking up the space into models of varying fidelity based on the environment, robot, and the task that defines the interaction between them. Every model contains a representation of state and action description to generate those next state variables. The action space might be differential equations of motion that include controls. For organization, models increase fidelity when they are supersets of other models. One criteria of this is if models can be transformed to be representations of others models [by either eliminating variables or holding them constant] then they tend to be higher fidelity. Higher fidelity models can be direct supersets of lower fidelity models by including more states in the space, or containing more control inputs. Higher fidelity models also include additional constraints that cause the model to be closer to the real world. The more correct a model is at representing the real world (such as modeling slip) or the larger the space that is considered in the model (higher dimensionality by considering dynamics, a superset) the higher fidelity the model is.

Our work assumes the models are correct, but on real robot platforms models are never perfect. Sensors are also not always accurate. Sometimes it is just not possible to have enough sensory information to detect a clear state for the robot. To have models more accurately reflect reality we believe we need to look at how uncertainty is added. There are examples of trying to account for uncertainty in the world that can not be captured or explained in even the most detailed model. This raises questions of how to account for uncertainty. A very conservative way may just increase padding around uncertain states. A less conservative way might apply uncertainty to specific dimensions of the state or parts of the action space. We would like to address uncertainty in the model graph and in the model selection process.

Uncertainty also expands smart model selection with the addition of partially observable data. The robot can now make probabilistic choices. It can have different strategies that trade-off being more conservative and safe versus getting to the goal in a faster riskier way. Incorporating risk attitudes captures the trade-off between variance and mean in the distribution of reward outcomes. The notion of risk can be considered during model selection. This tradeoff, between choosing a model that does or does not have uncertainty, is a function of how much risk the robot should take. If the robot chooses no uncertainty, it assumes a higher risk of failure, but with a possible higher reward payoff of a much shorter and faster path. The variance of the robot's utility becomes very large. It will either create a low cost path that accomplishes its task or it will fail and require recovery. If the robot does consider uncertainty, it acts much more risk-averse. Overall, it will tend to fail less often, but its expected path cost is higher.

Initial work was applied to a simple simulated navigation task. The task was for a simulated two wheeled differential drive robot, in Gazebo, to go between a start and end location. Success was based on reaching the goal, and failure was attributed to obstacle collision. The world included obstacles such as walls and overhangs, and a periodic dynamic opening and closing door. We planned using RRTs for fast plan generation that accommodate differential constraints. We hand generated the models, borrowing from previous motion planning models, based on the environment (2D or 3D), the robot chassis (x, y, z, theta), and the interaction of the robots motions. Motions expanded from straight line motion, to s-curve motion, and finally to s-curve motion with time and velocity constraints.

Our initial approach attempts to detect the parts of the lower fidelity plans that are infeasible for execution and repair them using re-planning through higher fidelity model selection. The model selection process uses prior plans to autonomously select the most applicable higher fidelity model in the hierarchy. This higher fidelity model is used to locally plan to an intermediate goal where the previous lower fidelity plan is resumed. This approach creates a mixed model plan. Initial results are consistent with demonstrating lower planning times without sacrificing execution success. For switching, while the average overall planning time (initial plus re-plans) is greater than the average time for the lowest fidelity plan alone, it is significantly less than the average for the higher fidelity spaces, while yielding success rates equal to the highest fidelity model space.