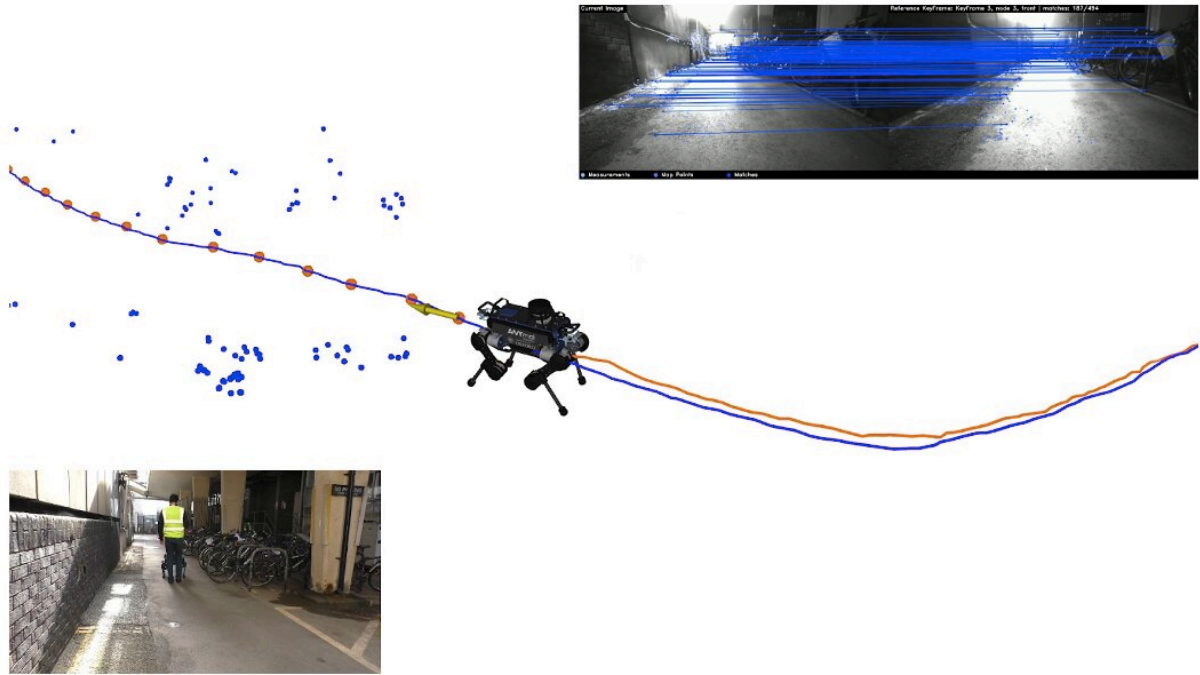# A locally reactive controller to enhance visual teach and repeat systems

February 8 2022, by Ingrid Fadelli

The researchers' original VT&R system only relied on the taught trajectory (in blue) and the visual input (top right) and had no notion of objects or people around it (bottom left). Credit: Mattamala, Ramezani, Camurri & Fallon.

To operate autonomously in a various unfamiliar settings and successfully complete missions, mobile robots should be able to adapt to changes in their surroundings. Visual teach and repeat (VT&R) systems are a promising class of approaches for training robots to adaptively

navigate environments.

As their name suggests, VT&R systems are based on two key stages: the teach and the repeat steps. During the teach step, the systems learn from demonstrations of paths taken by human operators. Subsequently, during the repeat stage, the robots try to replicate what the humans did in the demonstration, walking down the same path autonomously and as consistently as possible.

Researchers at the Oxford Robotics Institute have recently developed a new controller that could help to enhance VT&R systems. Their approach, presented in a paper published in *IEEE Robotics and Automation Letters*, could help to develop robots that are better at navigating unfamiliar environments.

"The recent paper is part of our work on VT&R navigation," Matias Mattamala, one of the authors, told TechXplore. "This is useful to quickly deploy robots to inspect new places and collect data without having to build a precise map of the environment. In our previous work, we demonstrated robustness to visual occlusions by switching between different cameras on the robot, such as when someone is walking by."
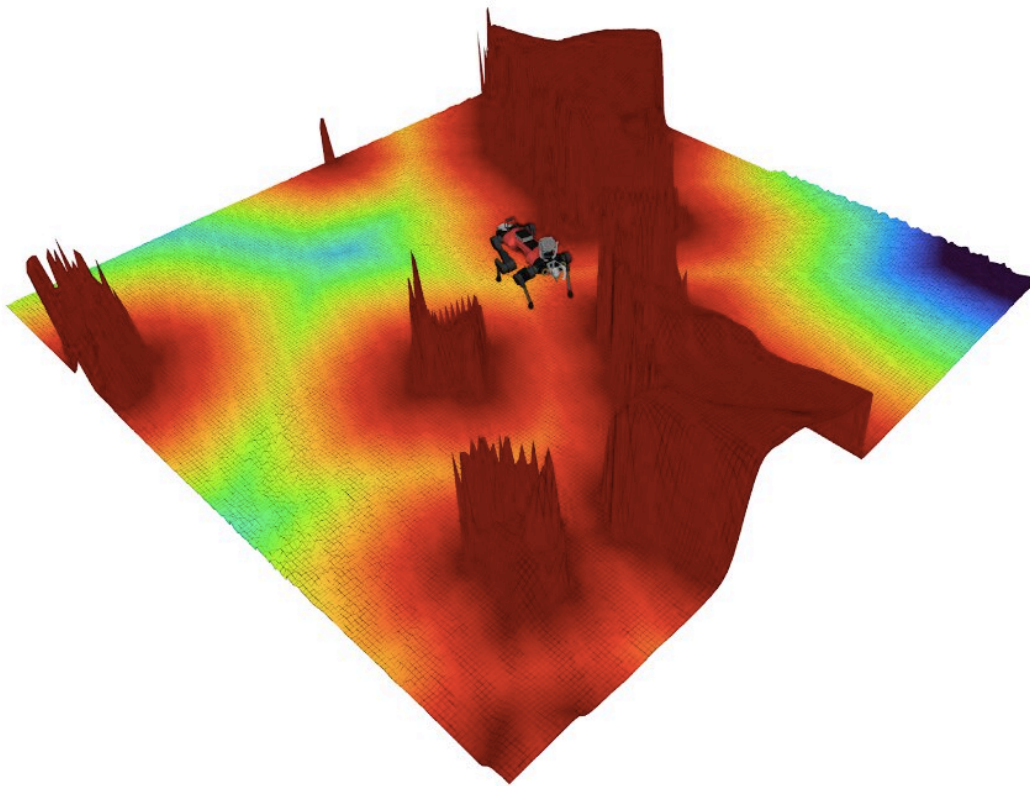
A Signed Distance Field (SDF) computed from the local map around the robot was used to generate repulsion forces to stay away from obstacles, especially in narrow corridors. Red means closer to obstacles and blue farther away. Credit: Mattamala, Chebrolu & Fallon.

In their previous studies, Mattamala and his colleagues were able to train models to access different cameras on a robot at different times, using data they collected during human demonstrations. Despite this remarkable achievement, their models did not allow robots to actively avoid potentially obstacles in their surroundings while replicating the trajectory demonstrated by human agents.

"We started to work on this 'safety layer' some time ago and our recent paper presents it fully functional," Mattamala explained. "Our controller

is based on a recent approach developed by Nvidia called [Riemannian Motion Policies](#) (RMP)."

The controller developed by the researchers partly resembles potential field controllers, tools that allow robots to compute a combination of different forces, such as a attraction forces (i.e., those driving them toward completing a goal) and repulsion forces (i.e., those helping them to stay away from obstacles), to ultimately determine what direction to move in. Nvidia's RMP approach, however, takes their controller one step further, as it introduces dynamic weights (called metrics) that leverage these forces in different ways, depending on the state of the robot.
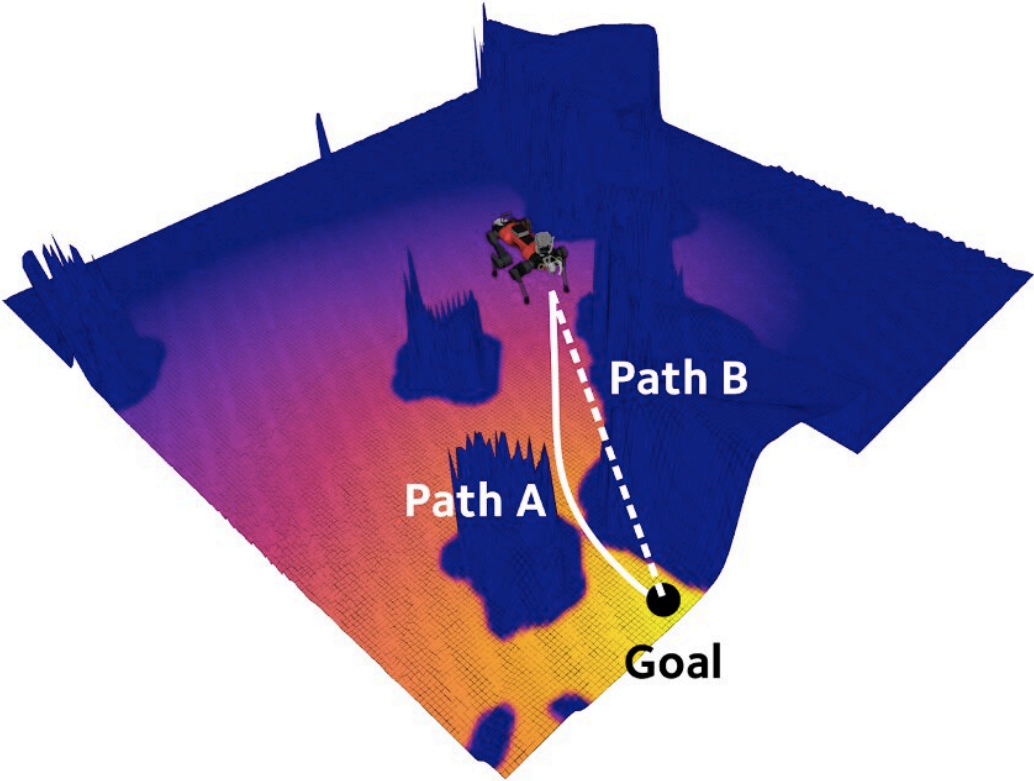
The Geodesic Distance Field (GDF) was used to determine directions (a

gradient) that helped the robot to reach a goal without trespassing obstacles. For example, instead of trying to follow a straight path to the goal (Path B) that is the easiest thing to do, the GDF will generate a gradient to follow Path A instead. Credit: Mattamala, Chebrolu & Fallon.

"For example, you don't need to always avoid obstacles, but only when you are close to them or pointing in their direction," Mattamala explained. "In this way, you can prevent some situations in which attraction and reaction forces cancel each other."

The interacting forces processed by the team's controller are computed from a local map that is generated on the fly and adapts as a robot moves in its surrounding environment. By analyzing this local map, the system can generate fields that are easy to interpret and can be used as data to enhance a robot's navigation skills. This includes a signed distance field (SDF), which characterizes obstacles, and a geodisc distance field (GDF), which conveys the closest distance to a goal or target location. When processing these fields, the controller accounts for the fact that there is a certain amount of space in the surrounding environment that the robot cannot move in or traverse.

"In our study, we were able to explore novel control techniques such as RMP, which have so far only been applied to robot manipulators or small wheeled robots," Mattamala said. "In addition, we deployed our controller on the [ANYbotics'](link) ANYmal quadruped and performed [closed-loop experiments in a decommissioned mine](link), which was quite exciting to test."

The team's experiment in a decommissioned mine near Wiltshire, UK, was a 60 m long path. The robot traversed it without collisions, despite the challenging lighting conditions and narrow passages. Credit: Mattamala, Chebrolu & Fallon. Picture by Oliver Barlett.

In contrast with other previously proposed approaches, the controller created by Mattamala and his colleagues is intrinsically reactive, as it does not require robots and developers to plan ahead and predict the obstacles a robot will encounter in a specific environment. Interestingly, in their evaluations, the team found that by using better environment representations to generate attraction and reaction forces, they could achieve similar results to those attained by models that plan for missions in advance.

"For example, we placed some obstacles blocking the reference path and the robot was able to go around without planning," Mattamala explained. "We also extended our VT&R system to work with fisheye cameras, such as the [Sevensense](link) Alphasense rig that we used in our experiments. We achieved comparable results to previous experiments with Realsense cameras, which demonstrated the flexibility of our system."

The team's experiment in a decommissioned mine near Wiltshire, UK, was a 60 m long path. The robot traversed it without collisions, despite the challenging lighting conditions and narrow passages. Credit: Mattamala, Chebrolu & Fallon. Picture by Oliver Barlett.

So far, the researchers have tested their controller in a series of indoor cluttered spaces and in an underground mine. In these initial experiments, their system achieved very promising results, suggesting that it could soon help to enhance the navigation capabilities of both existing and newly developed mobile robots. Notably, the controller could be applied to a variety of systems, as it only requires a local map generated using data collected by depth cameras or LiDAR technology.

Image 6: Some environments in which the team plan to test the robot's autonomous navigation in their future work. These are some examples of pictures taken while testing the RLOC walking controller, developed by Siddhant Gangapurwala. Credit: Gangapurwala. Picture by Oliver Barlett.

In their next studies, Mattamala and his colleagues plan to apply and test their controller on other robots developed in their lab. In addition, they would like to evaluate its performance in a broader range of dynamic, real-world environments.

"Our future work considers extending our VT&R system to achieve the long-term visual navigation of legged robots in industrial and natural environments," Mattamala explained. "This requires (1) better visual localization systems, since drastic appearance changes due to lighting or weather conditions will challenge our current system, and (2) better walking controllers to achieve reliable navigation in rough terrain, which should interact with the high-level navigation. Imagine teaching the robot to traverse forest trails or to hike along a mountain trail, and then repeating the trajectory autonomously, no matter the terrain or the weather—that's what we aim to achieve."