Twitter Thread by Florian Shkurti

Florian Shkurti
@florian_shkurti

Inspired by @michellearning I will be sharing 10 robotics papers that I really enjoyed reading (made me think and raised my heart rate) when I was a grad student. One paper per day. Starting tonight, ending on xmas day. Here we go: ■■ 1/n

Note: this is not necessarily an exhaustive list of the best or most influential papers in robotics, nor is it a list of my current favorite papers. These are papers that I found compelling and I still enjoy re-reading from time to time. The order is random. 2/n

[Paper #1] LQR-Trees by Russ Tedrake https://t.co/nYDaHYJryk. I love this paper. It's about feedback control and motion planning with probabilistic guarantees. Given a desired goal state, sample an initial state and find a nominal trajectory from the initial state to the goal 3/n

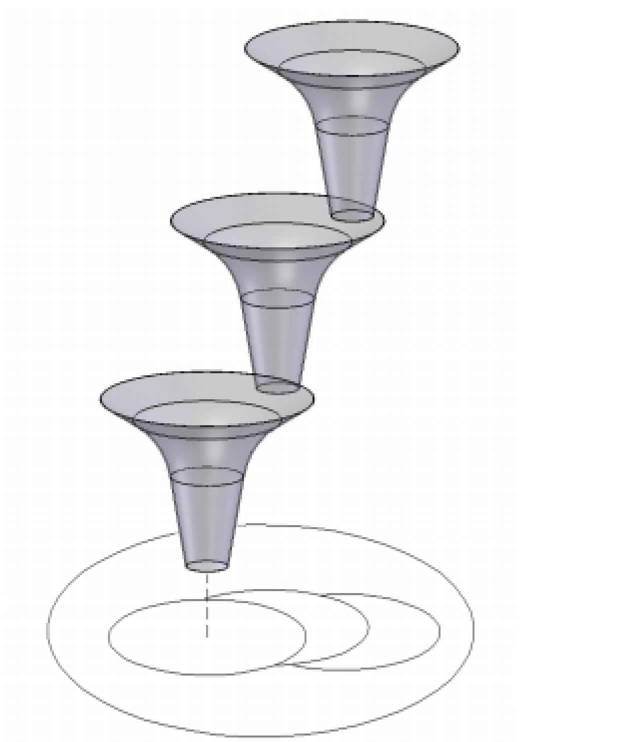

Fig. 1: Cartoon of motion planning with funnels in the spirit of [4].

using LQR control (linear dynamics, quadratic cost). So far, so good. Then it gets interesting: it sets out to explicitly compute the region of stability of the computed controller without doing sampling. It underestimates this region by trying to compute a Lyapunov function 4/n

that will certify stability in the region. It only considers Lyapunov functions that are polynomials of deg N and optimizes for a positive function using sums-of-squares programming. So, now you have a nominal trajectory to the goal and a region of stability around it. 5/n

Repeat this process backwards by setting the new goal to be in the computed region of stability, and you have a region of stability that grows over time and (a) covers space, (b) can lead to complex global behaviors from simple linear controllers that are locally valid, 6/n

and (c) is one way to address compositionality in control. So, does all of this work in practice? Yes, with a few asterisks: the major bottleneck is solving the sum-of-squares program that computes the region of stability. It has been shown to be practical for low-dim systems 7/n

This paper, published in ~2010, was influenced by Pablo Parrilo's PhD thesis on semidef programming (2000), as well as older ideas about funnels from Matt Mason https://t.co/4MxHCJQy7l (1985) Rob Burridge, Alfred Rizzi, and Dan Koditschek https://t.co/LjNnY0ItD3 (1999) 8/n

Some of the limitations of the LQR Trees paper were addressed by @Majumdar_Ani in his PhD thesis https://t.co/m5jcbmZ3x5 I encourage you to read it and check out his other papers https://t.co/vso7qQchGp 9/n