

Pros and Cons of truncated Gaussian EP in the context of Approximate Inference Control

Marc Toussaint, TU Berlin

mtoussai@cs.tu-berlin.de

October 29, 2009

Approximate Inference Control (AICO, [1]) is a method for solving Stochastic Optimal Control (SOC) problems. A core step in this method is to approximate the belief in each time step, that is, compute a distribution over the system state given the coupling to the previous state, the next state, and the current constraints or goals. Since constraints and goals may be arbitrarily non-linear (specified in arbitrary task spaces) exact inference is infeasible. The core motivation for introducing AICO was to leverage the variety of existing *approximate* inference methods to suggest new efficient approximations in the context of SOC. In [1] we used Gaussian message approximations which can be computed from local LQG approximations of the system dynamics and costs. We also mentioned the possibility to use Expectation Propagation and truncated Gaussians for inference under hard constraints and limits. With this abstract we discuss in more detail the approach, its promises and limitations, and report on experiments with AICO and truncated Gaussians for robotics problems.

Truncated Gaussians vs. collision potentials. Efficient motion optimization algorithms are either based on sampling (explicitly testing whether a configuration or state is feasible, RRTs, road maps, etc) or based on some optimization principle which assumes a smooth cost function (gradient-based, CHOMP, DDP/iLQG, etc). In the latter case, to deal with hard (non-continuous) constraints one has to introduce local approximations to the hard constraint. A typical example is a repelling squared potential: Given a list of object pairs with distances d_i below a margin m one may define a squared cost function $C_{\text{col}} = \alpha_{\text{col}} \sum_i (1 - d_i/m)^2$. Using the Jacobian J_i of pair distances one can pull this back to a squared cost function in joint space. Efficient optimal control methods like DDP or AICO can be used with this cost approximation. However, this approach implies that collision constraints are now in trade-off with other cost terms. Choosing the constant α_{col} becomes a hand-tuning problem. Also, fixing the margin parameter m is often undesired: during slow and controlled object manipulation the fingers may get very close to objects, during a fast large-scale reaching motion the margin should much higher.

Using truncated Gaussians seems to provide a principled solution to hard constraints – here are the pros of

truncated Gaussians:

- a) it directly corresponds to minimizing the probability of collision (maximizing the likelihood $P(\text{non-collision})$),
- b) its effect on a belief is variance dependent,
- c) it is parameter free.

To point a): In the probabilistic setting, avoiding collisions corresponds to introducing a binary random variable which indicates collisions and conditioning this variable to zero. This evidence results in a discontinuous message over the state space, which can be approximated by a heavy-side function $\theta(x)$ along a hyperplane (normal to the distance Jacobian J_i). Given a Gaussian belief $b'(x)$ in state space which accounts for all other information (fwd & bwd messages, squared costs) we use expectation propagation (EP) to compute a message which approximates the effect of the truncation: We compute the moments of the truncated belief $\theta(x)b'(x)$ which defines the Gaussian belief $b(x) \approx \theta(x)b'(x)$. The constraint evidence is finally represented as the Gaussian message $\mu(x) \propto b(x)/b'(x)$. To summarize, this procedure directly approximates the effect of conditioning on non-collision. It does not rely on inventing a heuristic cost function that indirectly might lead to avoiding collisions.

To point b): The effect of the truncation of the belief in state space very much depends on the variance of the belief. For instance, a high-variance belief around a configuration close to collision is truncated much more than a low-variance belief. This effect is particularly interesting seems desirable in the context of the full SOC framework.

To point c): There are no parameters implicit in this approach. In place of the margin parameter, the typical mean distance to collision depends on the variance of the belief. In place of the α_{col} parameter, the effect of the truncation depends on the likelihood of collision.

Examples. Figure 1 illustrates the effect of using truncated Gaussian EP in AICO for trajectory planning. Most interestingly: in the collision potential case we see that collisions have no effect when the belief mode (which is used to compute the local LQG approximation) is outside the proximity margin of $m = 0.1$, independent of the belief's variance. We also performed experiments on simulated 7DoF robotic configurations (reaching with

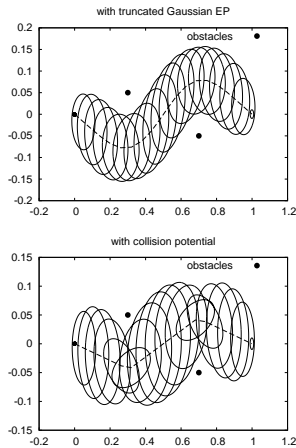


Figure 1: The result of AICO when using truncated Gaussian EP vs. squared collision potentials (margin $m = 0.1$) on a toy problem. Illustrated are the posterior beliefs of 20 time steps of a trajectory with random walk prior starting at $(0, 0)$ conditioned to reaching $(1, 0)$ and avoiding obstacles.

an arm), where the proximity distances were computed using a collision detection library (SWIFT++) and the truncating hyperplanes was transformed into joint space using the Jacobians of these distances. The additional computational cost of the truncation and EP are negligible in this specific setup. The resulting posterior trajectories successfully avoid collisions – typically by keeping rather large distances to the objects. We will report on these results on the poster.

Limitations in practice. However, when applying the method in such robotic scenarios we realized some limitations (which we will expand on more on the poster):

a) Every collision pair has to be handled by its own truncation: In the case of collision potentials, all (near) collisions are summarized by a single cost scalar C_{col} , thereby adding only a single task variable to the optimization problem. In the case of handling con-

straints with truncated Gaussians, every collision pair contributes a new constraint and thereby adds to the computational cost. In relevant scenarios (e.g., manipulation) there may be hundreds of such collision pairs which would lead to very high computational costs. We tried an alternative: Since we can in principle handle constraints in any task space we can impose a constraint on a single scalar that subsumes all collisions, e.g., C_{col} . However, we failed with this trick since the severe non-linearity of this scalar prohibited convergence of EP and AICO. We were only successful with translating each near-collision to a constraint.

b) Typical collision detection thresholds lead to discontinuities in the algorithm that may prevent convergence of EP and AICO: Typically, in simulations one discards pairs of objects to be collision relevant when their distance is above a certain threshold ϵ . In effect, the collision library discontinuously switches between reporting and not reporting a critical distance. This can have negative effects in the case of truncated Gaussian: In our experiments, when we set ϵ to naturally small values, the discontinuous switching of active constraints prohibits EP and thereby AICO to converge. The method works nice, if we set ϵ rather large – but this increases inefficiency.

In conclusion, EP with truncated Gaussians in the context of motion control is theoretically very elegant and straight-forward in the context of AICO. We are not yet sure, however, if the practical limitations mentioned can be solved. For further details on computing moments of truncated Gaussians and corresponding software please see <http://user.cs.tu-berlin.de/~mtoussai/source-code/>.

- [1] M. Toussaint. Robot trajectory optimization using approximate inference. In *Proc. of the 26rd Int. Conf. on Machine Learning (ICML 2009)*, 2009.