RESTRICTING MOTION ALONG TANGENT DIRECTIONS IN BUNDLE ADJUSTMENT

ADAM WILLIAMS

1. Context

Consider a bundle adustment problem optimizing a least-squares cost function f over a variable x in a manifold M:

$$f(x) = r(x)^\top W r(x).$$

Here, $r : M \to \mathbb{R}^d$ is the residual and W is a weighting matrix (probably an inverse covariance). Assume that we will refine the cost function by taking a Gauss-Newton step:

$$\begin{aligned} &\delta = (J^\top W J)^{-1} J^\top W r \\ &x \leftarrow x \oplus \delta \end{aligned}$$

where $J : T_x M \to \mathbb{R}^d$ is the Jacobian of $r(x)$ at x:

$$r(x \oplus \delta) \approx r(x) + J\delta.$$

where $\oplus$ is an exponential-like mapping from $T_x M$ to M.

In certain situations, it is useful to restrict the parameter motion at a given step to a strict linear subspace of $T_x M$. Let $\mathcal{S} \subset T_x M$ be such a subspace, and let

$$\Pi : T_x M \to \mathcal{S}$$

be the projection of the tangent space onto $\mathcal{S}$. This writeup describes a convenient formulation of the Gauss-Newton step in the restricted tangent space.

2. Method

Let $\Pi^\dagger$ be a right inverse of Π, so that $\Pi\Pi^\dagger$ is the identity on $\mathcal{S}$. The adapted Gauss-Newton step works as follows:

(1) Replace $J : T_x M \to \mathbb{R}^d$ with $J\Pi^\dagger : \mathcal{S} \to \mathbb{R}^d$.
(2) Compute the Gauss-Newton step $\delta^* \in \mathcal{S}$ with the modifed Jacobian.
(3) Lift the subspace step and apply as usual:

$$\begin{aligned} \delta &= \Pi^\dagger \delta^* \\ x &\leftarrow x \oplus \delta. \end{aligned}$$

3. Applications

The simplest (and primary motivating) use case for the above method is to restrict motion to (the span of) a specified set of coordinate axes in the tangent space T_xM. In particular, if T_xM has axes $(e_1, \dots, e_m)$, and we wish to restrict motion along $e_{i_1}, \dots, e_{i_k}$, then Π is *the* $(m-k) \times m$ *matrix obtained by removing rows* $i_1, \dots, i_k$ *from the* $m \times m$ *identity matrix.* And $\Pi^\dagger = \Pi^\top$.

3.1. **Example: Holding the position of a sensor fixed.** Consider a sensor whose extrinsics are parameterized as a group element $g = \begin{bmatrix} R & t \end{bmatrix} \in \mathrm{SE}(3)$, representing the "sensor-from-body" transformation, acting on body-frame points via $g \cdot p = Rp + t$. Assume a local parameterization via exponential coordinates $\delta = \begin{bmatrix} \omega \\ u \end{bmatrix} \in \mathfrak{se}(3)$, so that Jacobians are computed with respect to the sensor frame, and updates are applied in the sensor frame:

$$g \leftarrow \exp(\delta) g.$$

Note that, in this parameterization, the position of the sensor in the the body frame is given by

$$g^{-1} \cdot 0 = -R^\top t.$$

To restrict the motion such that the position of the sensor in the body frame remains fixed (without modifying or replacing the underlying exponential parameterization), we simply project onto the "rotation" axes:

$$\Pi = \begin{bmatrix} I_{3\times 3} & 0_{3\times 3} \end{bmatrix}.$$

The modified Gauss-Newton step using $J\Pi^\top$ thus computes the optimal step restricted to the "rotation-only" subspace of $\mathfrak{se}(3)$, which gets lifted into $\mathfrak{se}(3)$, resulting in an update of the form $\delta = \begin{bmatrix} \omega \\ 0 \end{bmatrix} \in \mathfrak{se}(3)$. Hence the left-update applied to g is of the form $\begin{bmatrix} S & 0 \end{bmatrix}$, and the updated value of g is

$$g \leftarrow \exp(\delta) g = \begin{bmatrix} SR & St \end{bmatrix}.$$

Observe that the sensor position in the body frame has not changed:

$$g^{-1} \cdot 0 = -(SR)^\top St = -R^\top S^\top St = -R^\top t.$$

3.2. **Example: Parameterizing unit vectors as motion-restricted rotations.** Consider a variable n taking values in the unit sphere S^2. Given an existing left-update exponential parameterization $R \leftarrow \exp(\omega) R$ of the rotation group SO(3), the motion restriction method provides an easy parameterization of such a sphere point.

We obtain this parameterization by noting that a point n on the sphere can be identified with the subspace of SO(3) that rotates, say, the $+X$ axis onto n, and that the extra degree of freedom in SO(3) corresponds to rotation about n. In fancier mathematical language, we are noting that $S^2 \cong \mathrm{SO}(3)/\mathrm{SO}(2)$.

Hence, we can model a unit vector as an element of SO(3), as long as, at each step, we restrict rotation around the unit vector that the rotation represents. In particular, we represent n as any rotation matrix

$$R = \begin{bmatrix} n^\top \\ b^\top \\ c^\top \end{bmatrix}$$

whose inverse maps $\begin{bmatrix}1\\0\\0\end{bmatrix}$ onto n, and use the subspace projection

$$\Pi = \begin{bmatrix}0 & 1 & 0\\0 & 0 & 1\end{bmatrix}.$$

Then the modified Gauss-Newton step computes an update in the "roll-free" rotation subspace, having the form $\delta = \begin{bmatrix}0\\\omega_y\\\omega_z\end{bmatrix}$. This update represents a small rotation that leaves the $+X$ axis fixed at the infinitesimal level, and thus provides a minimal parameterization of the tangent space $T_n S^2$ at the current parameters.