Supplementary Materials
for
"Dynamic Time-of-Flight"

Michael Schober
Max Planck Institute for
Intelligent Systems

Amit Adam* Omer Yair* Shai Mazor*

Sebastian Nowozin
Microsoft Research

Abstract

Supplementary materials to the main paper, containing further technical details and additional results.

1. Live Video Demonstration

We submit a video with audio commentary (`video-2685.mp4`) showing further results on live scenes. The video is annotated and self explanatory and all depth results shown are without any spatial filtering. We summarize the main points this video demonstrates:

- Robustness of the approach: we show the results on four varied sequences which include both camera and object motions.
- Realtime performance: the results were obtained by running the trained regression trees as described in Section 4 in the main paper, achieving 30 frames per second.
- The dynamic TOF model reduces depth noise: we demonstrate this in static image and 3D point cloud visualizations.
- The method handles motion seamlessly—any remaining (minor) artifacts are due to the regression tree approximation. Because the inference function being approximated is deterministic the tree approximation can be made more accurate using more training samples, multiple trees, or deeper trees.

2. Motion Map Statistic

We compute the *motion map* statistic for each pixel as

$$M = 1 - \mathbb{E}_{\vec{\theta}^{(1:2)}}\left[\frac{\omega P(\vec{\theta}^{(2)})}{\omega P(\vec{\theta}^{(2)}) + (1-\omega) Q(\vec{\theta}^{(2)} \mid \vec{\theta}^{(1)})}\right], \tag{1}$$

*A. Adam, O. Yair, and S. Mazor were with Microsoft AIT, Haifa, Israel. {`email.amitadam`, `omeryair`, `smazor.shai`}`@gmail.com`.

where the expectation in (1) is over the posterior beliefs, $P(\vec{\theta}^{(1:2)}|\vec{R}^{(1:2)})$.

3. MCMC Inference Accuracy

Markov chain Monte Carlo (MCMC) is an approximate inference method that generates samples from the exact posterior distribution only asymptotically [1]. Therefore, in practice we need to truncate the Markov chain after a finite number of steps and the amount of steps required depends on both the difficulty of the posterior distribution and the *mixing time* of the chain, that is, how fast the Markov chain visits all regions of the posterior distribution.

The choice of number of steps can be made either statically or adaptively while running the chain. To perform adaptive truncation we could use popular convergence diagnostics [1], however, maintaining the diagnostics brings about additional runtime overheads, so we adopt the simpler static strategy of using a fixed number of MCMC iterations.

In the experiments of the main paper we adopt the choice of 2^{17} iterations, because this offers a reasonable tradeoff between accuracy and speed. To confirm this choice, we perform a simple validation experiment as follows.

We generate 4096 samples from the prior and a fixed number m of MCMC iterations for burn-in, not recording any statistics. Then we again use m iterations, this time collecting the posterior mean statistic using an efficient running mean algorithm. We evaluate $m \in \{2^{13}, 2^{14}, \dots, 2^{19}\}$. The runtime is linear in m and we show accuracy results in Figure 1, evaluated using the known ground truth.

The results confirm that after 2^{17} iterations there are only small additional gains.

4. Parameter Constraints

In Section 3.4 of the main paper we discuss the MCMC procedure and mention that the parameters are constrained to their valid range. We now provide additional details on how we achieve this during inference.

There are two parts to constraining the parameters: *first*, how to avoid leaving the feasible set during MCMC pertur-

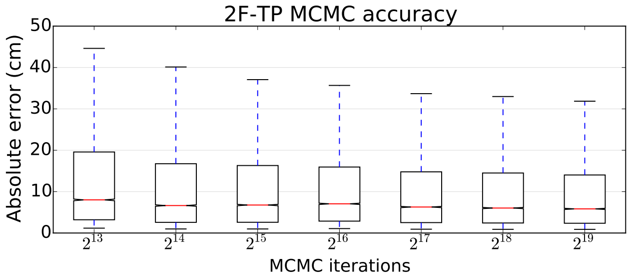

Figure 1: Evaluation of the 2F-TP inference accuracy as a function of the MCMC iterations used. The box plots show the 25/50/70 percentiles of the absolute depth errors in cm, the whiskers show the 10/90 percentiles.

bations, and *second*, how to evaluate the correct likelihood function under constraints.

To remain in the feasible set we can simply reject any MCMC proposal move that would violate a parameter constraint. The resulting Markov chain preserves the correct stationary distribution [2].

To evaluate the likelihood function correctly, we need to truncate and renormalize the density functions affected by truncation. The per-frame observation model uses the uniform distribution and we can achieve truncation by suitably shrinking the interval. The dynamic motion model is more involved, as it involves the non-uniform Laplace distribution either as an additive perturbation (for $t^{(s+1)}$ given $t^{(s)}$ and likewise for $t_2^{(s+1)}$), or as a multiplicative factor (for $\rho^{(s+1)}$ given $\rho^{(s)}$, and $\lambda^{(s+1)}$ given $\lambda^{(s)}$).

For example, to evaluate the likelihood of the motion model $P(\theta^{(s+1)}|\theta^{(s)})$ one term is based on the ratio $\rho^{(s+1)}/\rho^{(s)}$. For this term the truncated likelihood for the multiplicative factor $f_\rho = \rho^{(s+1)}/\rho^{(s)}$ (Equation (13) in the main paper) yields a truncated Laplace distribution given as

$$\text{TruncLap}(\rho^{(s+1)}/\rho^{(s)}, \mu = 1, b_\rho, \ell = \rho_{\text{min}}/\rho^{(s)}, u = \rho_{\text{max}}/\rho^{(s)}),$$

where $[\ell, u]$ is the truncation interval. Effectively the interval only allows factors f_ρ that ensure $\rho^{(s+1)} \in [\rho_{\text{min}}, \rho_{\text{max}}]$. For the other parameters similar constraints apply, all leading to a truncated Laplace distribution. We evaluate all transition likelihoods through the apropriately truncated Laplace distribution; likewise, when sampling from the motion model prior we also sample from the correctly truncated Laplace distribution.

We now give a summary of the truncated Laplace distribution. Although straightforward to derive, we are not aware of any prior description of the distribution.

4.1. Truncated Laplace Distribution

We consider the case when a Laplace distribution is conditioned on the interval $[\ell, u]$, effectively assigning zero probability outside this interval. The resulting distribution has a density function that is the same as the Laplace density but requires a suitable renormalization. We derived the truncated Laplace density and give the detailed form of the truncated Laplace density in Figure 2.

4.2. Sampling from the Truncated Laplace Distribution

We now give a procedure how to sample from $\text{TruncLap}(\mu, b, \ell, u)$. In the main paper this is needed whenever a sample from the prior is generated, such when initializing the MCMC procedure and when attempting a Metropolized independence sampling transition.

Because we must have $\ell < u$ there are only three possibilities to consider:

1. If $\mu \geq \ell$ and $\mu \leq u$.

   We perform simple rejection sampling as follows. We generate samples from the non-truncated Laplace distribution using the exponential reduction: if $x \sim \text{Exp}(1/b)$, and $y \sim \text{Exp}(1/b)$, then $z = x - y + \mu$ is a sample from $\text{Laplace}(\mu, b)$. If $z \notin [\ell, u]$ we repeat the procedure. While this procedure could be inefficient if $u - l \ll b$ we have not observed any problems in practice.

2. If $\mu < \ell$.

   We use inverse transform sampling by first generating $U \sim \text{Uniform}([0,1])$, then computing

   $$z = \mu - b \log\left(\exp(\log(1-z) - \tfrac{\ell}{b}) + \exp(\log(z) - \tfrac{u}{b})\right). \tag{3}$$

   It can be shown that z is distributed as $\text{TruncLap}(\mu, b, \ell, u)$.

3. If $\mu > u$.

   By symmetry we reduce this case to the previous one: we first generate $z' \sim \text{TruncLap}(\mu, b, 2\mu - u, 2\mu - \ell)$, then compute $z = 2\mu - z'$. Again z is correctly distributed.

References

[1] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.

[2] J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, 2001.

$$
\mathrm{TruncLap}(x;\mu,b,\ell,u) = \begin{cases} \frac{\exp(-|x-\mu|/b)}{2b-b\exp((\ell-\mu)/b)-b\exp(-(u-\mu)/b)}, & \text{if } \ell \le \mu \le u, \\ \frac{\exp(-|x-\mu|/b)}{b\exp(-(\ell-\mu)/b)-b\exp(-(u-\mu)/b)}, & \text{if } \mu < \ell, \mu < u, \\ \frac{\exp(-|x-\mu|/b)}{b\exp((u-\mu)/b)-b\exp((\ell-\mu)/b)}, & \text{if } \ell < \mu, u < \mu, \\ 0, & \text{if } x \notin [\ell, u]. \end{cases} \tag{2}
$$

Figure 2: Probability density function for the truncated Laplace distribution $\mathrm{TruncLap}(\cdot;\mu,b,\ell,u)$ with mean μ, dispersion parameter b and bound interval $[\ell,u]$.