N-Dimensional Gaussians for Fitting of High Dimensional Functions: Supplementary Material

Stavros Diolatzis
stavros.diolatzis@intel.com
Intel Labs
France

Tobias Zirr
tobias.zirr@intel.com
Intel Labs
Germany

Alexander Kuznetsov
alexander.kuznetsov@intel.com
Intel Labs
USA

Georgios Kopanas
georgios.kopanas@inria.fr
Inria, Université Côte d'Azur
France

Anton Kaplanyan
anton.kaplanyan@intel.com
Intel Labs
USA

1 APPLICATION INPUT PARAMETERIZATION

We include here the details of the input parameterization for the two applications which also defines the dimensionality of the Gaussian mixtures.

1.1 Global Illumination with Variability

In this application we parameterize the Gaussians over typical inputs are usually given to a neural network for shading a synthetic scene. As mentioned in the main text the dimensions are: world position xyz, view direction xyz_{dir}, albedo rgb and roughness r. World position is normalized to be between -1 and 1 using the scene bounding box. View direction is not encoded in any way it just includes the xyz components of the normalized view direction vector. As albedo we use the reflectance values of any diffuse or rough materials. Finally, roughness is also not encoded in any way, and it is parameterized as roughness^2 or a of the GGX distribution.

1.2 Volumetric Radiance Fields

For the 6D scenario (world position xyz, view direction xyz_{dir}), the view direction is computed as the vector from the camera to the world position mean $m_x yz$ of each Gaussian. This is the same parameterization as 3DGS which is used to evaluate the spherical harmonics encoding. In our case we don't encode the view direction in any way, it just serves as 3 extra dimensions for our Gaussian mixture.

2 PROJECTION FROM 6D TO 3D

To utilize the splatting approach from [Kerbl et al. 2023] we project our 6D Gaussians to 3D by conditioning them on the current camera view direction. Given the 6D mean $\mathbf{m}$ and covariance $\mathbf{V}$ we partition their elements into:

$$\mathbf{m} = \begin{bmatrix} \mathbf{m}_1 \\ \mathbf{m}_2 \end{bmatrix}, \quad \mathbf{V} = \begin{bmatrix} \mathbf{V}_{11} & \mathbf{V}_{12} \\ \mathbf{V}_{21} & \mathbf{V}_{22} \end{bmatrix} \tag{1}$$

where m_1 is the 3D world position mean, and m_2 is the xyz of the viewing direction. $\mathbf{V}_{11}$ is the 3D covariance of the world position, $\mathbf{V}_{22}$ is the 3D covariance of the viewing direction and $\mathbf{V}_{12}, \mathbf{V}_{21}$ are the cross-covariances. To condition on a given viewing direction $\mathbf{c}$ the conditional 3D mean $\mathbf{m}_{\text{cond}}$ and 3D covariance $\mathbf{V}_{\text{cond}}$ is:

$$\mathbf{m}_{\text{cond}} = \mathbf{m}_1 + \mathbf{V}_{12}\mathbf{V}_{22}^{-1}(\mathbf{d} - \mathbf{m}_2) \tag{2}$$

$$\mathbf{V}_{\text{cond}} = \mathbf{V}_{11} - \mathbf{V}_{12}\mathbf{V}_{22}^{-1}\mathbf{V}_{21} \tag{3}$$

Since we are dealing with unnormalized Gaussians, the value $G_{\mathbf{V}_{\text{cond}}}$ of the derived 3D Gaussian has to be adjusted by $G_{\mathbf{V}}(x)$ where $\mathbf{x} = [\mathbf{m}_{\text{cond}}, \mathbf{c}]$. That is to say the maximum value the conditional 3D Gaussian can have is the value of the 6D Gaussian at the point of the new mean given the condition.

REFERENCES

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.* 42, 4 (2023).