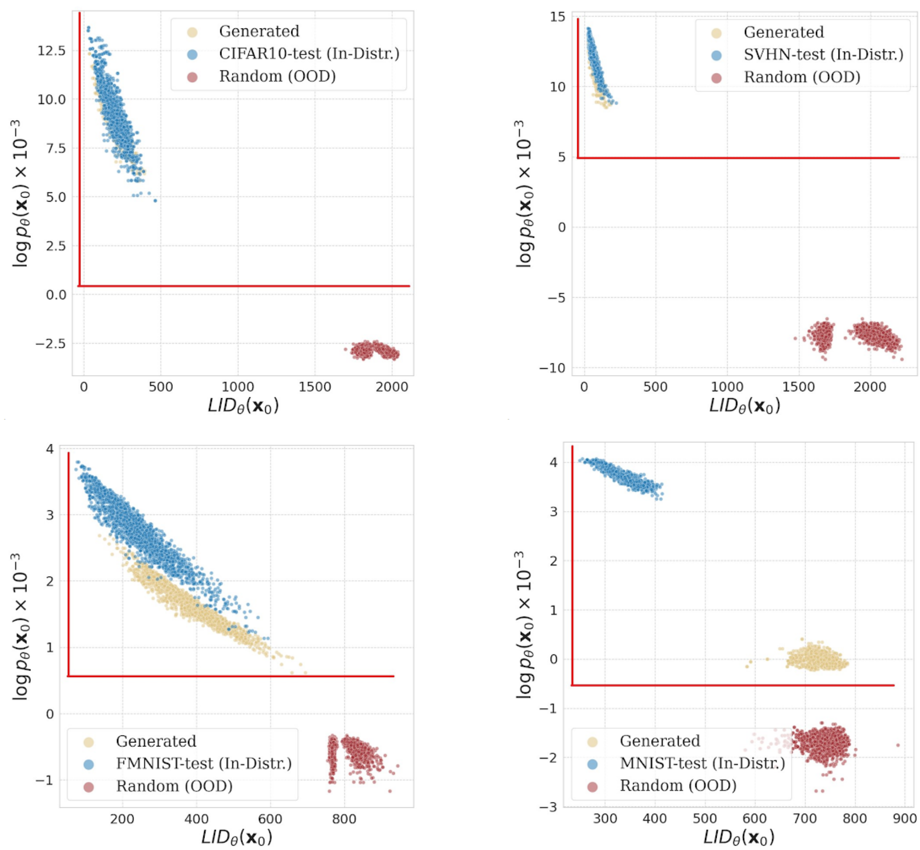

Figure 1: Density vs. LID estimates for flows trained on various datasets against random noise.

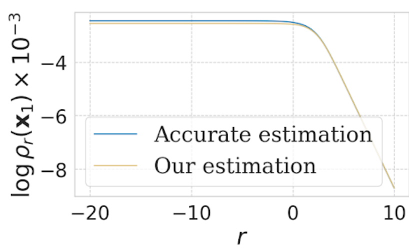

(a) Convolution of the first mode

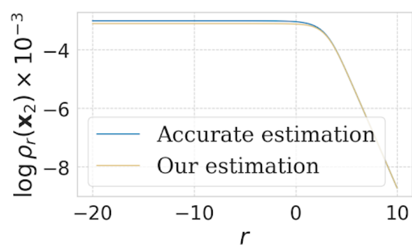

(b) Convolution of the second mode

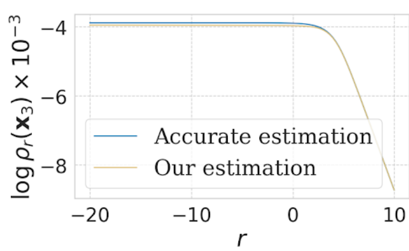

(c) Convolution of the third mode

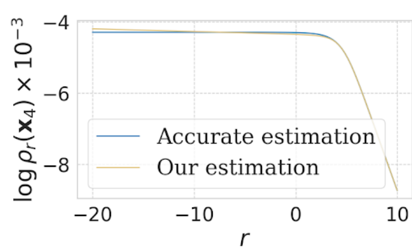

(d) Convolution of the fourth mode

Figure 2: The estimated convolution $\log \rho_r(\mathbf{x}_i)$ for different modes of the Gaussian mixture.

1 On the Quality of the Linear Approximation

In this section, we mathematically assess the accuracy of the approximation to ρ_r, as discussed in subsection 4.1, by establishing error bounds. To advance our discussion, we first introduce two key mathematical operators essential for deriving these bounds.

Definition 1.1. For positive scalars σ and δ, the *constrained Gaussian convolution operator* $\phi_\sigma^\delta : \mathcal{F} \to \mathcal{F}$ on the family of smooth functions $\mathcal{F}$ mapping from $\mathbb{R}^d$ to $\mathbb{R}^+$ takes in an arbitrary function h and outputs g as follows:

$$\phi_\sigma^\delta(h) = g \text{ s.t. } g(\mathbf{x}) \coloneqq \int_{B_{\sigma\delta}(\mathbf{x})} h(\mathbf{x} - \mathbf{u}) \cdot \mathcal{N}(\mathbf{u}; \mathbf{x}, \sigma^2 \cdot \mathbf{I}) \cdot d\mathbf{u}. \tag{1}$$

Here, the integral is on the ℓ_2 ball of radius $\sigma \cdot \delta$. The *unconstrained Gaussian convolution operator* is also defined similarly with the exception that the integration is over the complement of the ball:

$$\bar{\phi}_\sigma^\delta(h) = g \text{ s.t. } g(\mathbf{x}) \coloneqq \int_{\mathbb{R}^d \setminus B_{\sigma\delta}(\mathbf{x})} h(\mathbf{x} - \mathbf{u}) \cdot \mathcal{N}(\mathbf{u}; \mathbf{x}, \sigma^2 \cdot \mathbf{I}) \cdot d\mathbf{u}. \tag{2}$$

Note that $\phi_\sigma^\delta + \bar{\phi}_\sigma^\delta$ results in the normal convolution operator, and when the input of the operator is the density function p_θ, we have that $\rho_r(\mathbf{x}_0) = \phi_\sigma^\delta(p_\theta)(\mathbf{x}_0) + \bar{\phi}_\sigma^\delta(p_\theta)(\mathbf{x}_0)$ for $\sigma = e^r$ and any δ. Now we will provide an upper bound for the unconstrained convolution operator which will help us provide a global error margin for $\hat{\rho}_r(\mathbf{x}_0)$:

Lemma 1.1. *Given a bounded function $h : \mathbb{R}^d \to \mathbb{R}^+$ where $M \coloneqq \sup_{\mathbf{x}\in\mathbb{R}^d}(h(\mathbf{x}))$, we have that*

$$\bar{\phi}_\sigma^\delta(h)(\mathbf{x}_0) \leq M \cdot \left[(\delta/d)e^{1-(\delta/d)}\right]^{d/2}, \tag{3}$$

when $\delta > d$.

Proof.

$$\begin{aligned}
\bar{\phi}_\sigma^\delta(h)(\mathbf{x}_0) &= \int_{\mathbb{R}^d \setminus B_{\sigma\delta}(\mathbf{x}_0)} h(\mathbf{x}_0 - \mathbf{u}) \cdot \mathcal{N}(\mathbf{u}; \mathbf{x}_0, \sigma^2 \cdot \mathbf{I}) \cdot d\mathbf{u} \\
&\leq \int_{\mathbb{R}^d \setminus B_{\delta}(\mathbf{x}_0)} M \cdot \mathcal{N}(\mathbf{v}; \mathbf{0}, \mathbf{I}) \cdot d\mathbf{v} \qquad \text{Change of variables } \mathbf{v} \coloneqq (\mathbf{u} - \mathbf{x}_0) \cdot \sigma^{-1} \\
&= M \cdot P(\chi_d^2 > \delta^2) \leq M \cdot \left[(\delta/d)e^{1-(\delta/d)}\right]^{d/2},
\end{aligned}$$

where χ_d^2 denotes a Chi-squared random variable with d degrees of freedom. The last inequality is a well-known Chernoff bound on the survival function of the Chi-squared distribution. □

Given a flow matching function f_θ which is diffeomorphic, a high-quality local approximation to f_θ around $\mathbf{x}_0$ would result in a high-quality approximation of p_θ via the change of variables formula. Now we will present error bounds for our core estimator $\hat{\rho}_r(\mathbf{x}_0)$.

Theorem 1.2. *Assume R specifies a region in which an approximation $\hat{p}_\theta$ is accurate up to a small error margin: $\forall \mathbf{x} \in B_R(\mathbf{x}_0) : |\hat{p}_\theta(\mathbf{x}) - p_\theta(\mathbf{x})| \leq e(R)$. If $M' := \sup_{\mathbf{x} \in supp(p_\theta)}(|p_\theta(\mathbf{x}) - \hat{p}_\theta(\mathbf{x})|)$ is finite and $e^r = \sigma < R/d$, then the total error between $\rho_r(\mathbf{x}_0)$ and $\hat{\rho}_r(\mathbf{x}_0)$ is bounded as follows:*

$$|\rho_r(\mathbf{x}_0) - \hat{\rho}_r(\mathbf{x}_0)| \leq e(R) + M' \cdot \left[\frac{R}{d\sigma} e^{1-\frac{R}{d\sigma}}\right]^{d/2} \tag{4}$$

Proof.

$$|\rho_r(\mathbf{x}_0) - \hat{\rho}_r(\mathbf{x}_0)| \leq \phi_\sigma^{R/\sigma}(|p_\theta - \hat{p}_\theta|)(\mathbf{x}_0) + \bar{\phi}_\sigma^{R/\sigma}(|p_\theta - \hat{p}_\theta|)(\mathbf{x}_0)$$
$$\leq \int_{B_R(\mathbf{x}_0)} |p_\theta(\mathbf{x}_0 - \mathbf{u}) - \hat{p}_\theta(\mathbf{x}_0 - \mathbf{u})| \cdot \mathcal{N}(\mathbf{u}; \mathbf{x}_0, \sigma^2 \cdot \mathbf{I}) \cdot d\mathbf{u} + M' \cdot \left[(R/d\sigma) e^{1-(R/d\sigma)}\right]^{d/2}$$
$$\leq e(R) + M' \cdot \left[(R/d\sigma) e^{1-(R/d\sigma)}\right]^{d/2}$$

□

Therefore, for a sufficiently small σ (translating to a sufficiently negative r) the total error bound of $\hat{\rho}_r$ is as small as the bound $e(R)$ obtained from the linearization. This analysis demonstrates that even though $\rho_r(\mathbf{x}_0)$ concerns taking a convolution over the entire support of p_θ, a high-quality local approximation of the density directly yields a high-quality approximation $\hat{\rho}_r(\mathbf{x}_0)$.