The Base Measure Problem and its Solution

Alexey Radul, Boris Alexeev {axch,alexeev}@google.com PROBPROG 2020, October 22-24

Example: What is the density after stretching a distribution on the circle?

x , y ~ uniform_on_unit_circle
x’, y’ = 2x, 20y

p(x’, y’) = ?

Easy, right?

- p(x, y) = $1/2\pi$ when $x^2 + y^2 = 1$
- Let f(x, y) = (2x, 20y)
- $|\det J_f| = 40$ everywhere
- Ergo p(x’, y’) = $1/80\pi$ when $(x'/2)^2 + (y'/20)^2 = 1$

Wrong!

Perimeter of ellipse $(x'/2)^2 + (y'/20)^2 = 1$ is about 81.28, much less than 80π.

The distribution isn’t uniform!

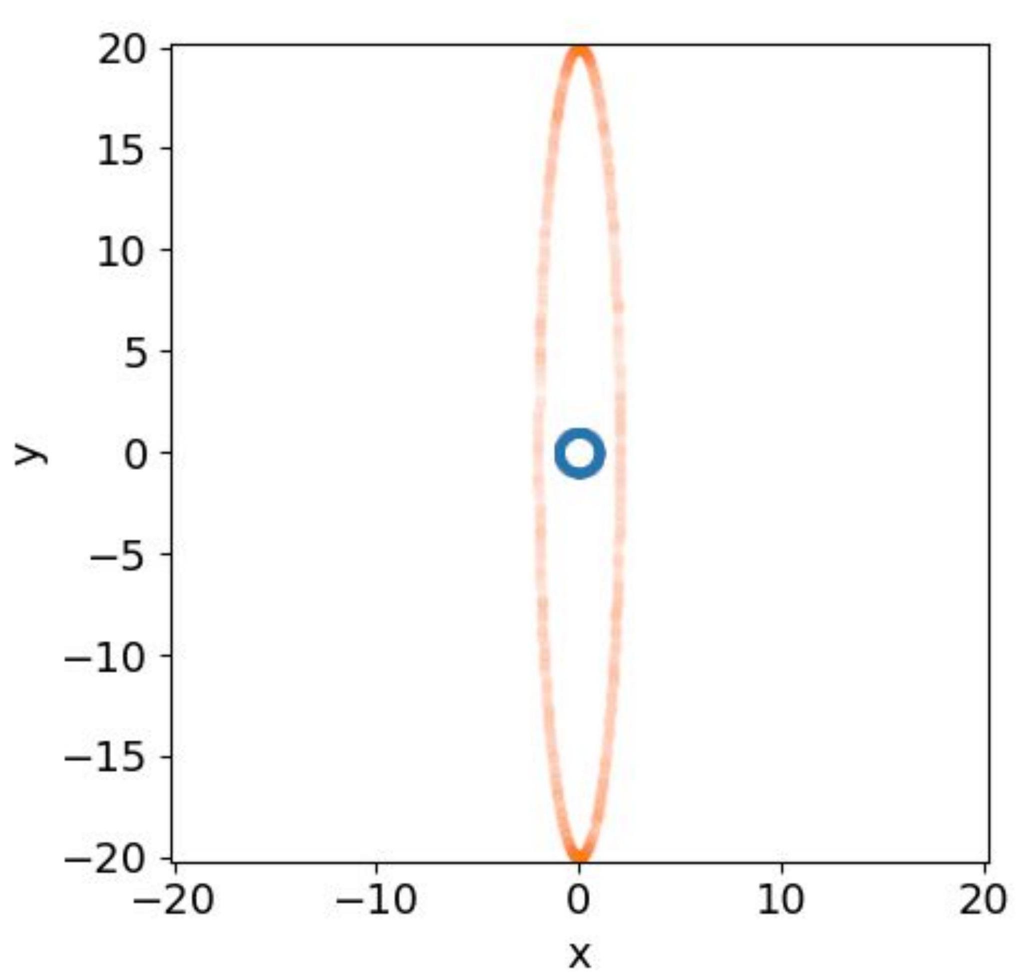

3000 samples of x’, y’ in orange.
3000 samples from the uniform distribution on the unit circle in blue for comparison.

Right answer

The density $1/2\pi$ is with respect to Lebesgue measure on the *circle*, not all of $\mathbb{R}^2$.

- Circle’s tangent at (x, y) is (-y, x).
- Directional derivative of f is (-2y, 20x).
- Change of arclength is $\mathrm{sqrt}(4y^2 + 400x^2)$.
- p(x’, y’) = $1/2\pi\mathrm{sqrt}(100x'^2 + y'^2/100)$.

In general:

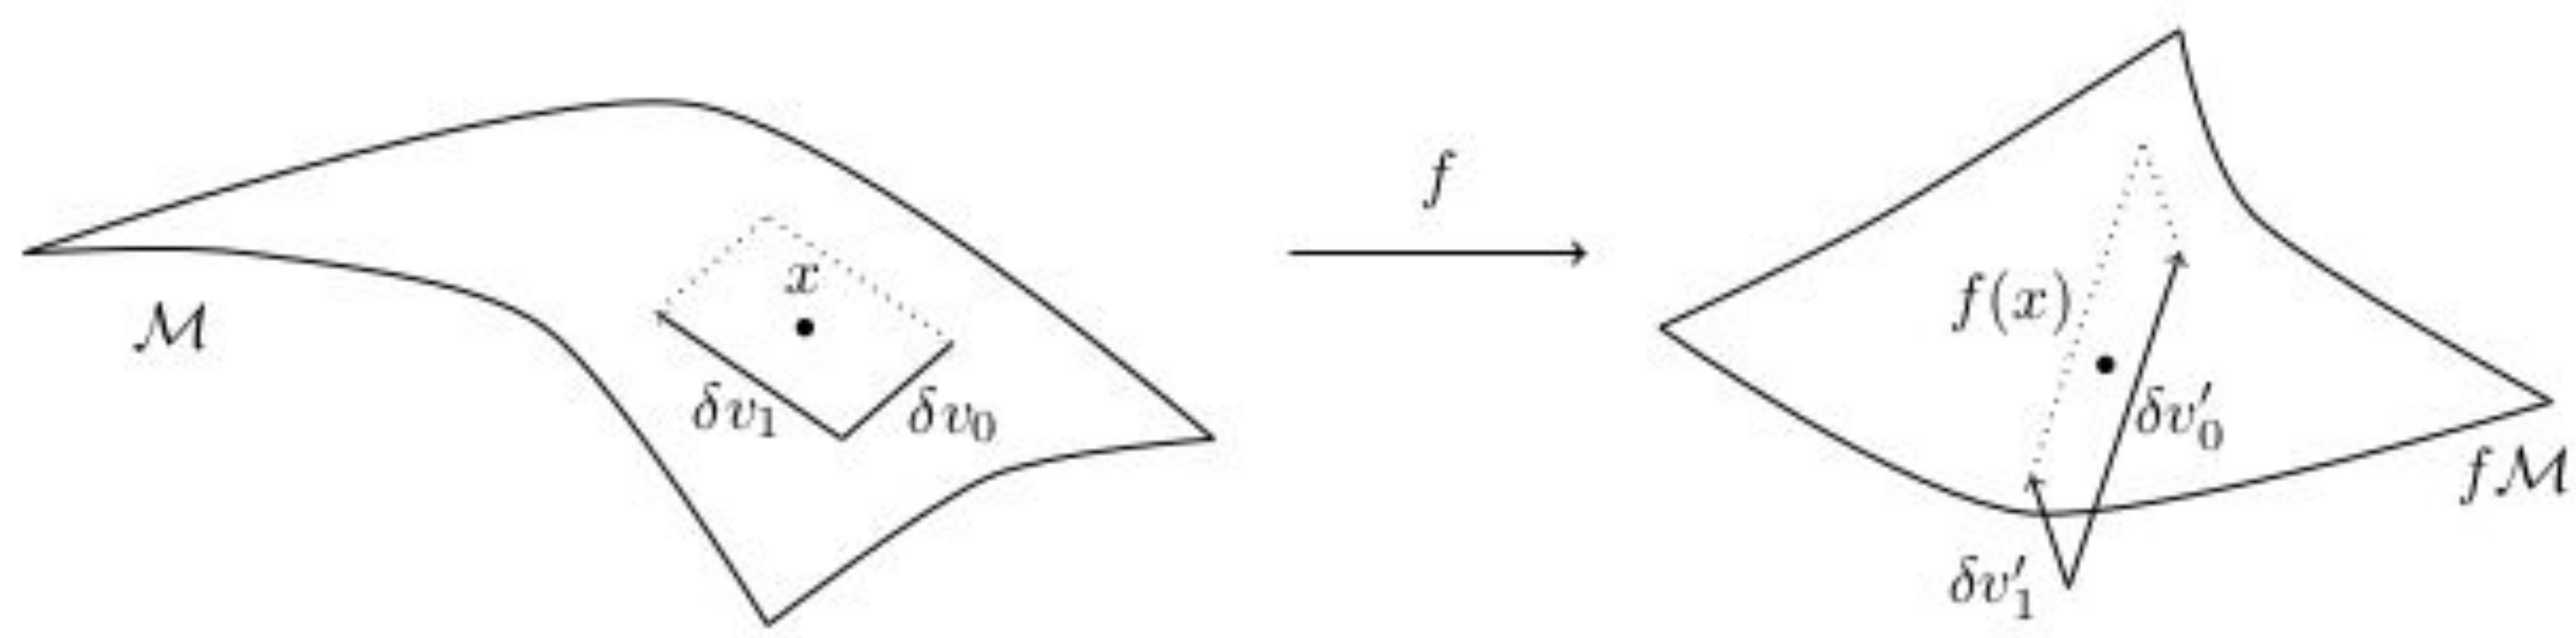

$p(x') = p(x)\ \mathrm{sqrt}(\det(VV^T)/\det(V'V'^T))$
where v_i is an arbitrary basis for the tangent space, and v_i' are those directional derivatives of f

When does this happen?

Whenever the base measure matters and is not Lebesgue on $\mathbb{R}^n$.

- Transforming discrete distributions embedded in $\mathbb{R}^n$.
- Transforming distributions on symmetric matrices, simplexes, spheres, etc.
- Reversible-jump MCMC on any of the above.
- MCMC or SMC with discrete + continuous observation model (e.g., Indian GPA problem).

Computation:

Log Jacobian determinants of bijections not enough.

Explicitly represent tangent space of support.

Automatic differentiation to compute directional derivatives.

Two-argument dispatch or Visitor pattern to cover efficient special cases.