

# Accurate Computation of the Logarithm of Modified Bessel Functions on GPUs

**Andreas L. Plesner**<sup>1</sup>, Hans Henrik Brandenborg Sørensen<sup>2</sup>, Søren Hauberg<sup>2</sup>

<sup>1</sup>ETH Zurich, [aplesner@ethz.ch](mailto:aplesner@ethz.ch)

<sup>2</sup>Technical University of Denmark, {hhbs, sohau}@dtu.dk

# We want to compute $\log I_\nu(x)$

For  $x = 0.5$  and  $\nu = 129.25$

- SciPy: -681.6595192185806
- MATLAB: -681.6595192185806
- boost: -681.6595192185805
- GNU Scientific Library (GSL): -681.6595192185805
- Mathematica: -681.6595192185806
- Wolfram|Alpha: -681.6595192185805

# We want to compute $\log I_\nu(x)$

For  $x = 0.5$  and  $\nu = 139.25$

- SciPy: -Infinity
- MATLAB: -Infinity
- boost: -744.4400719213812
- GNU Scientific Library (GSL): -Infinity
- Mathematica: -744.4400719213812
- Wolfram|Alpha: -744.5544364434038


# A brief introduction to (modified) Bessel functions

Bessel's Modified Differential Equation

$$x^2 \frac{d^2 y}{dx^2} + \frac{dy}{dx} - (x^2 + v^2)y = 0$$

$$\begin{aligned} & I_v(x) && K_v(x) \\ &= \left(\frac{1}{2}z\right)^v \sum_{k=0}^{\infty} \frac{\left(\frac{z}{4}\right)^k}{k! \Gamma(v+k+1)} \end{aligned}$$


Von Mises-Fisher Distribution




$$f(x; \mu, \kappa) = \frac{\kappa^{p/2-1}}{(2\pi)^{p/2} I_{p/2-1}(\kappa)} \exp(\kappa \mu^\top x)$$

# A brief introduction to (modified) Bessel functions

Modified Bessel Function of the First Kind



Modified Bessel Function of the Second Kind


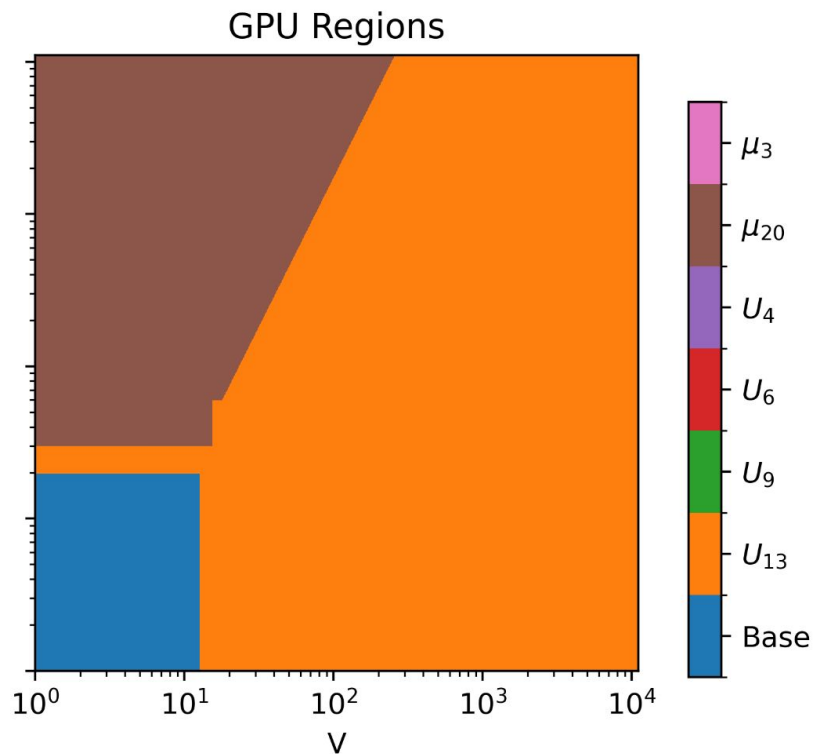
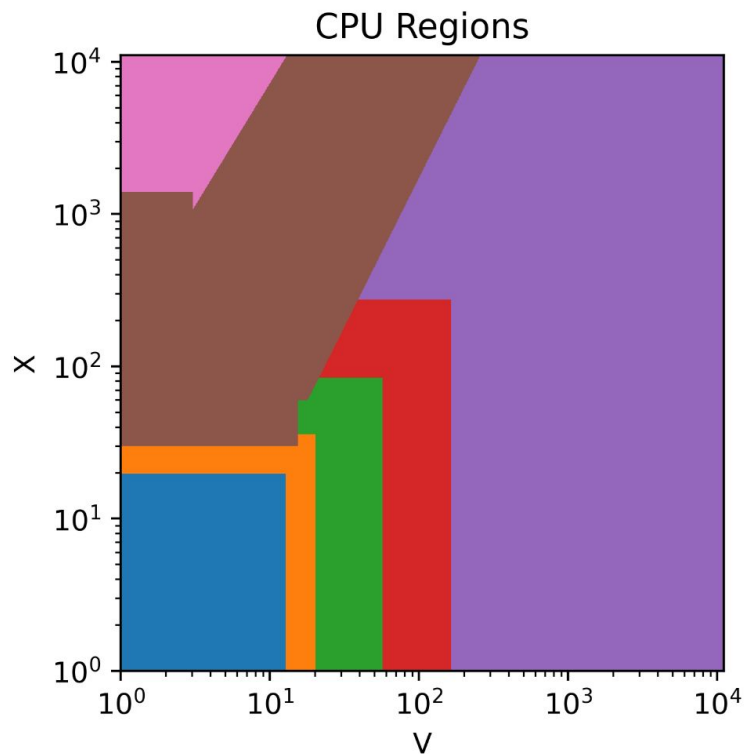


# What do these libraries do and how can we do better?

- Scaling the Bessel functions with exponentials
- Do intermediate computations in log scale
- Use asymptotic expressions
- Sort input values to prevent thread divergence on GPUs

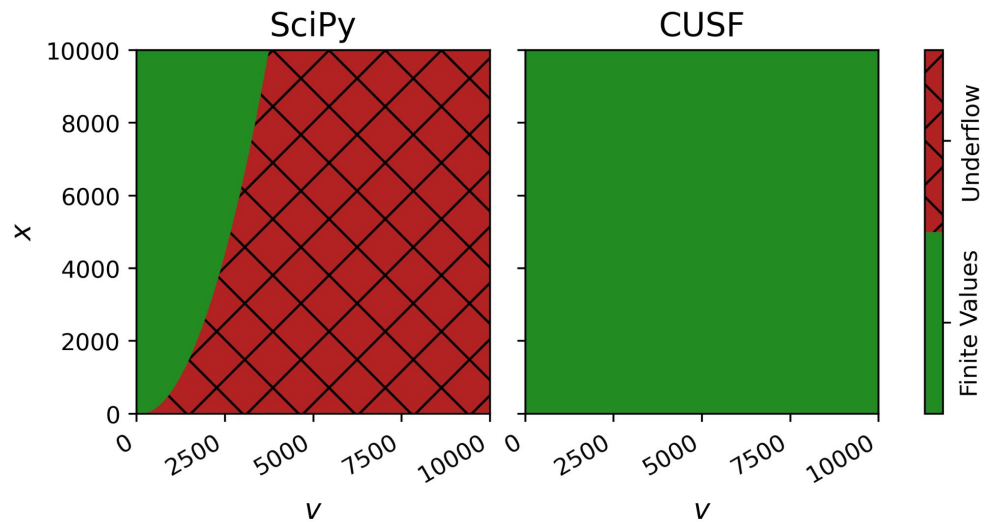
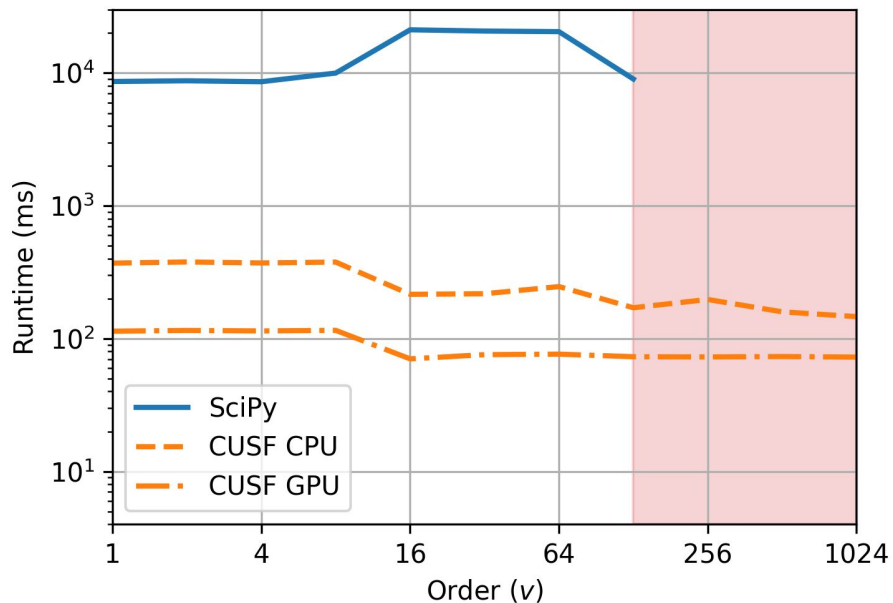
# Regions of the Different Expressions

$$I_v(x) \approx \frac{\exp(x)}{\sqrt{2\pi x}} \left( 1 + \sum_{k=1}^{\infty} (-1)^k \frac{\prod_{j=1}^k (\mu - (2j-1)^2)}{k!(8x)^k} \right), \mu = 4v^2$$



# How do we compare against SciPy for $\log I_\nu(x)$

**Median speedups of 35x and 77x for CPU and GPU, resp.**  
**Maximum speedups of 98x and 300x for CPU and GPU, resp.**



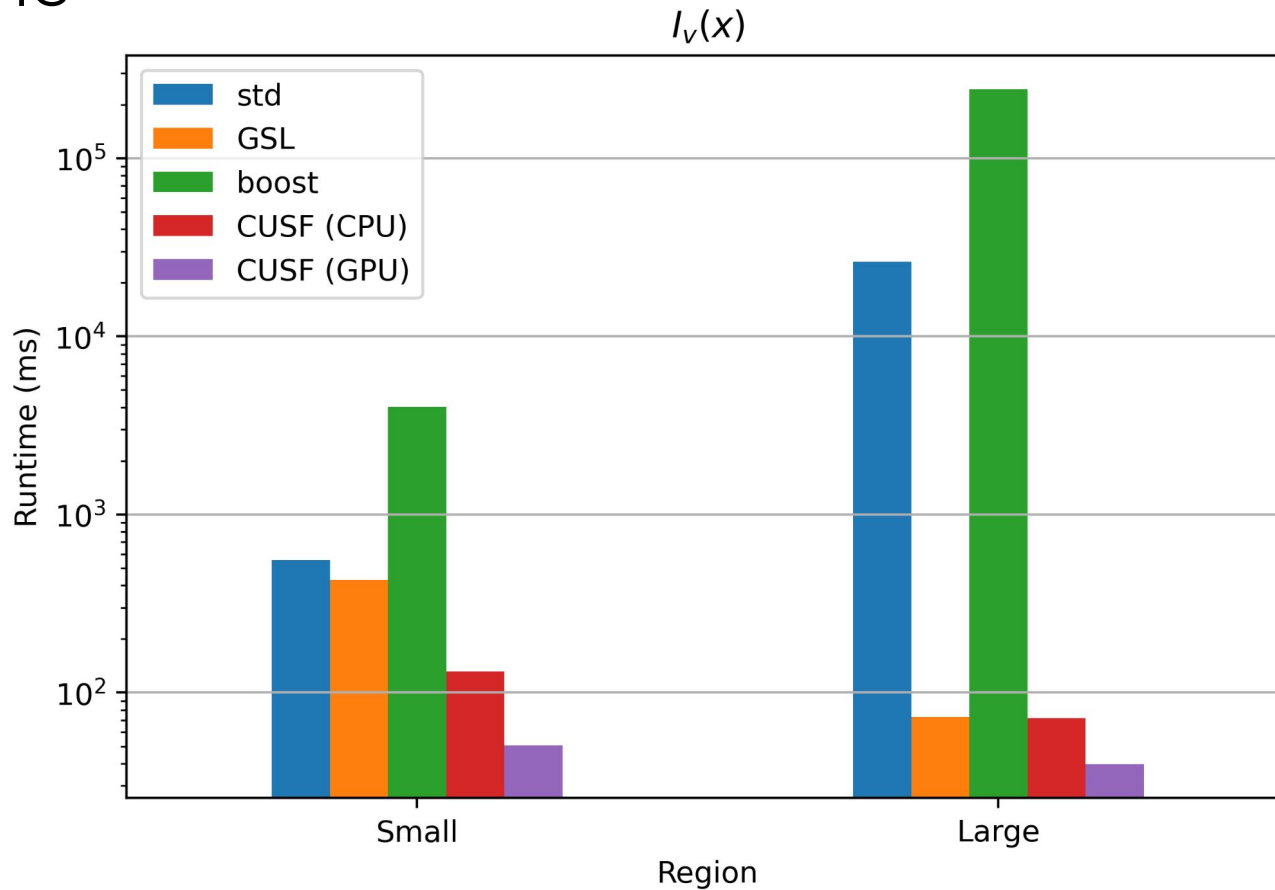


# Accuracy and robustness

Function	Region	Metric	std	GSL	Library Boost	CUSF (CPU)	CUSF (GPU)
$\log I_\nu(x)$	Small	Robustness	100%	99.98%	<b>100%</b>	100%	100%
		Median	$4.04 \times 10^{-16}$	$1.34 \times 10^{-16}$	0.0	$2.12 \times 10^{-16}$	$2.08 \times 10^{-16}$
		Maximum	$2.77 \times 10^{-6}$	$2.03 \times 10^{-7}$	$4.10 \times 10^{-8}$	$8.30 \times 10^{-4}$	$8.30 \times 10^{-4}$
	Large	Robustness	0.50%	44.13%	1.98%	<b>100%</b>	<b>100%</b>
		Median	$1.20 \times 10^{-16}$	0.00	0.00	$2.40 \times 10^{-16}$	$2.28 \times 10^{-16}$
		Maximum	$1.20 \times 10^{-5}$	$1.46 \times 10^{-15}$	0.00	$2.98 \times 10^{-13}$	$2.98 \times 10^{-13}$

Median speedups of 17x and 45x for CPU and GPU, resp.  
Maximum speedups of 3403x and 6150x for CPU and GPU, resp.

# Runtime



# Using CUSF with PyTorch

Imports

```
import numpy as np
import torch
from scipy.special import ive as scipy_ive
from cusf.bessel import iv_log as cusf_iv_log
```

Helper  
function

```
def scipy_iv_log(nu, z):
    np_result = np.log(scipy_ive(nu.cpu().numpy(), z.cpu().numpy())) + z.cpu().numpy()
    return torch.from_numpy(np.asarray(np_result)).to(device=z.device)
```

Data  
loading

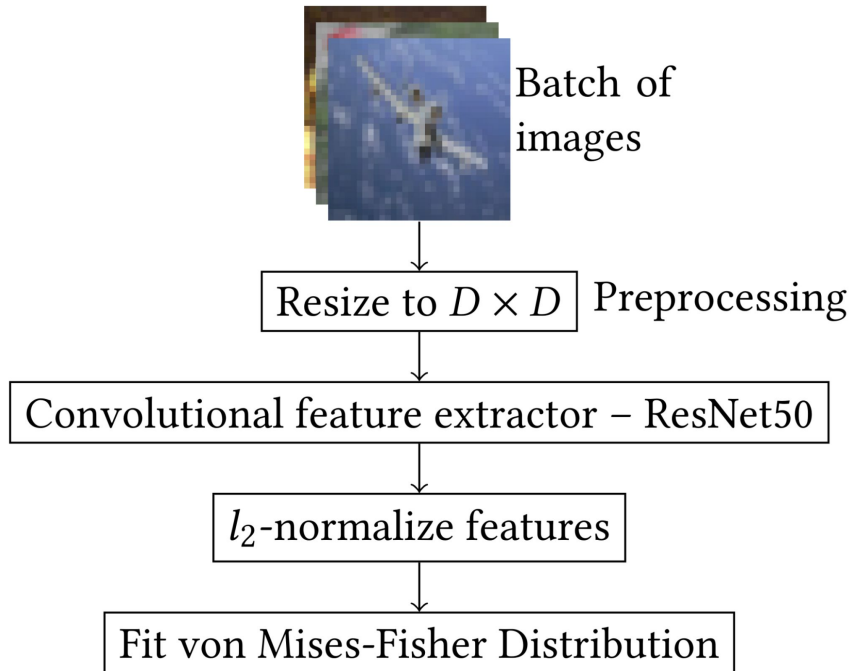
```
# Generate some data
torch.manual_seed(0)
device = torch.device(type='cuda')
N = 2000
M = 2000
z = torch.rand(N, M, device=device) * 150
nu = torch.rand_like(z) * 150
```

Doing  
stuff

```
# Evaluate
scipy_res = scipy_iv_log(nu=nu, z=z)
cusf_res = cusf_iv_log(nu=nu, z=z)
```

# Estimating von Mises-Fisher

## Pipeline



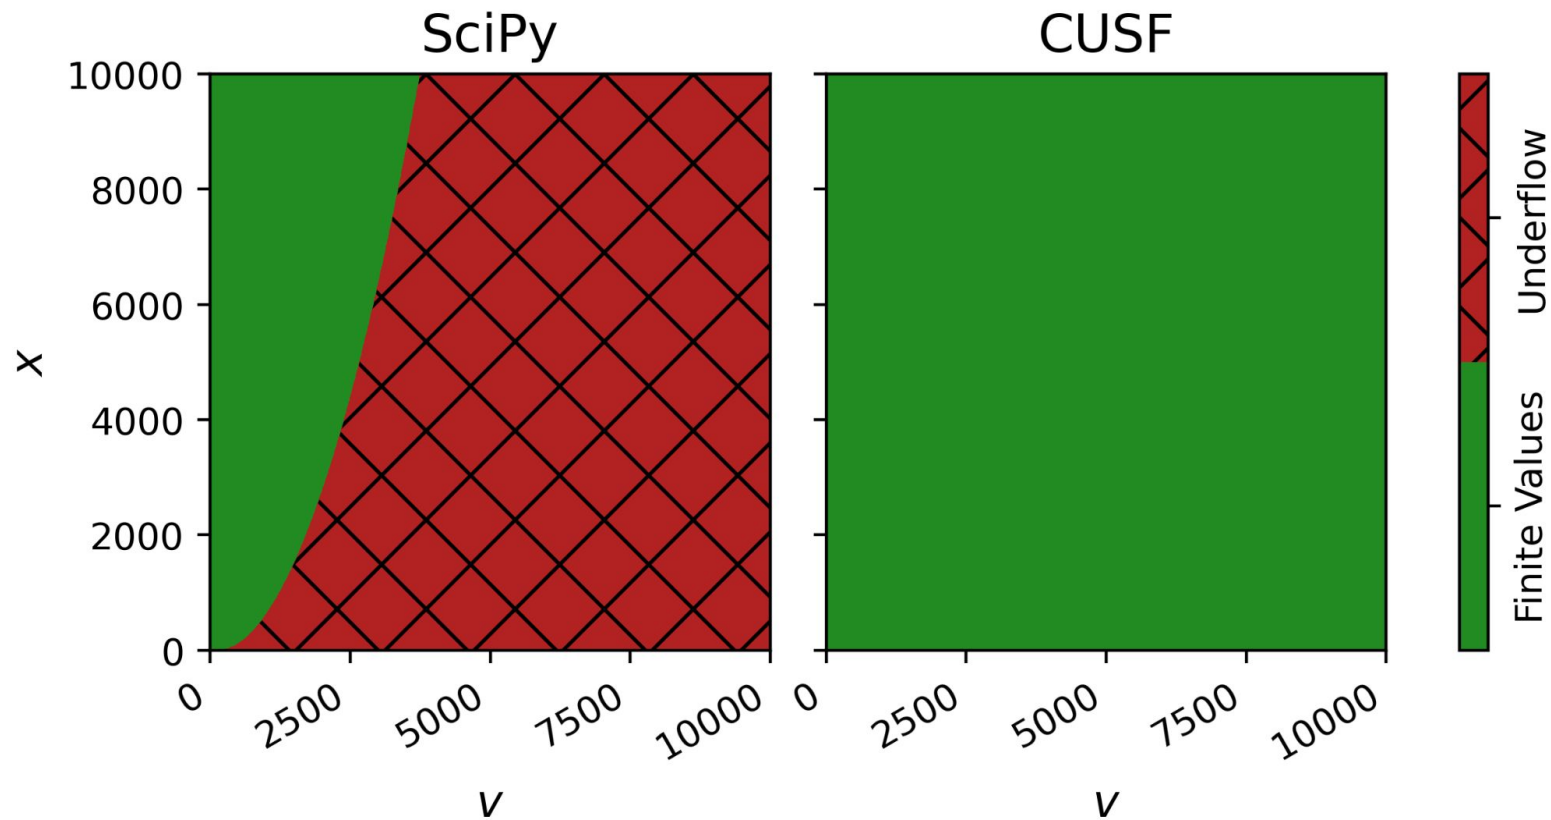
## Results

# of features	2048	8192	32768
Gradient free estimate	298.9091	1577.135	6681.31
Gradient estimate	298.9098	1577.405	6668.07
$\hat{\kappa}_0$	298.9127	1577.412	6668.08
$\hat{\kappa}_1$	298.9098	1577.405	6668.07
$\hat{\kappa}_2$	298.9098	1577.405	6668.07

# Is Mathematica the right benchmark?

Function	Region	Metric	Library				
			std	GSL	Boost	CUSF (CPU)	CUSF (GPU)
$\log I_\nu(x)$	Selected values	Robustness	100%	0%	100%	<b>100%</b>	<b>100%</b>
		Median	$2.28 \cdot 10^{-5}$	N/A	$2.28 \cdot 10^{-5}$	$1.53 \cdot 10^{-16}$	$1.53 \cdot 10^{-16}$
		Maximum	$8.30 \cdot 10^{-4}$	N/A	$8.30 \cdot 10^{-4}$	$3.07 \cdot 10^{-16}$	$3.07 \cdot 10^{-16}$

# Conclusion



Questions?

Andreas L. Plesner

[aplesner@ethz.ch](mailto:aplesner@ethz.ch)

# Comparison to CUDA

## Precision

Function	Region	Metric	Library					
			std	GSL	Boost	CUDA	CUSF (CPU)	CUSF (GPU)
$\log I_0(x)$	Small	Robustness	<b>100%</b>	100%	100%	100%	100%	100%
		Median	0.00	0.00	0.00	0.00	0.00	0.00
		Maximum	$1.11 \cdot 10^{-14}$	$3.08 \cdot 10^{-9}$	$9.07 \cdot 10^{-10}$	$9.07 \cdot 10^{-10}$	$3.68 \cdot 10^{-13}$	$3.68 \cdot 10^{-13}$
	Large	Robustness	61%	<b>100%</b>	61%	61%	100%	100%
		Median	0.00	0.00	0.00	0.00	0.00	0.00
		Maximum	$1.71 \cdot 10^{-16}$	$2.16 \cdot 10^{-16}$	$2.15 \cdot 10^{-16}$	$2.15 \cdot 10^{-16}$	$2.22 \cdot 10^{-16}$	$2.22 \cdot 10^{-16}$

## Speed

Function	Region	Library					
		std	GSL	Boost	CUDA	CUSF (CPU)	CUSF (GPU)
$\log I_0(x)$	Small	$3512.69 \pm 6.57$	$915.87 \pm 10.33$	$903.93 \pm 26.38$	<b><math>61.56 \pm 1.88</math></b>	$1665.29 \pm 5.96$	$433.00 \pm 18.63$
	Large	$19560.23 \pm 14.13$	$929.41 \pm 24.00$	$2080.96 \pm 33.07$	<b><math>50.15 \pm 0.26</math></b>	$301.89 \pm 25.13$	$187.14 \pm 21.57$