

# Understanding and Resolving NumPy's "RuntimeWarning: invalid value encountered in double\_scalars"

Authored by  
**Mohammed looti**

November 2, 2025

## RECOMMENDED CITATION

Mohammed looti (2025). *Understanding and Resolving NumPy's "RuntimeWarning: invalid value encountered in double\_scalars"*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8150>

For developers, data scientists, and computational engineers relying on high-performance numerical libraries like [NumPy](#) within the Python ecosystem, encountering numerical instability is an inevitable part of the job. One of the most common and critical signals of such instability is the appearance of a specific [RuntimeWarning](#). This warning is often misunderstood, but it flags a fundamental problem: the underlying floating-point arithmetic operation failed to produce a representable result, potentially leading to inaccurate model output or corrupted data processing.

The issue typically manifests when the results of intermediate calculations become either too small (underflow) or too large (overflow) for standard 64-bit precision to handle. This inability to represent the true mathematical value forces the system to output a special, indeterminate flag. The exact warning message that indicates this critical failure is standardized across many scientific Python environments:

### **runtimewarning: invalid value encountered in double\_scalars**

This warning is highly specific. It means that an operation involving 64-bit floating-point numbers--the standard precision often referred to as [double\\_scalars](#)--resulted in an invalid value. An invalid value usually means the calculation yielded a mathematically undefined result, such as  $0/0$  or  $\infty - \infty$ . When this occurs, Python assigns the result the special value [NaN](#) (Not a Number) or  $\infty$  (Infinity). Relying on results containing [NaN](#) severely compromises the integrity of any subsequent analysis or model training.

Addressing the root cause of this warning is non-trivial, as it requires moving beyond simple syntax fixes and employing robust mathematical techniques designed specifically for numerical stability. This comprehensive guide will dissect the mechanism behind the ``double_scalars`` warning and introduce the essential Log-Sum-Exp trick, demonstrating how to implement a stable and reliable fix using the powerful tools available in the SciPy library.

## **Understanding Floating-Point Precision and the IEEE 754 Standard**

To truly grasp why the ``double_scalars`` warning occurs, one must appreciate the limitations imposed by floating-point arithmetic. Most modern computing systems, including Python's numerical ecosystem powered by [NumPy](#), adhere to the **IEEE 754 standard** for floating-point representation. This standard dictates that a fixed number of bits--64 bits for standard double precision--are used to store a number's sign, exponent, and mantissa. While 64 bits offers enormous range and precision, it is finite, leading to inherent limitations when dealing with extremely large or extremely small numbers.

The core mechanism that triggers the [RuntimeWarning](#) is numerical **underflow or overflow**, particularly when interacting with transcendental functions like the exponential function ( $e^x$ ).

Numerical underflow happens when the result of a calculation is a non-zero number so close to zero that it cannot be accurately represented by the available exponent bits. The system is then forced to round this value down to absolute zero. Conversely, numerical overflow occurs when a result exceeds the maximum representable magnitude (approximately  $10^{308}$  for 64-bit floats), forcing the result to positive or negative infinity.

In the context of the ``double_scalars`` warning, the problem often arises when subsequent calculations attempt to use these compromised results. For instance, if an initial calculation underflows to zero, and that zero is then used as a denominator in a division operation, the resulting numerical overflow yields infinity. A more insidious scenario, and the one most relevant to this specific warning, involves the indeterminate form  $0/0$ . If both the numerator and denominator underflow to zero independently, their ratio is undefined, resulting in the special [NaN](#) value, which is precisely what the system flags as an invalid operation involving [double\\_scalars](#).

## The Specific Challenge: Ratios of Exponentials

The calculation pattern most prone to triggering the ``runtimewarning: invalid value encountered in double_scalars`` involves **ratios of sums of exponentials**. This structure is ubiquitous in fields such as machine learning, where it forms the basis of critical algorithms like the Softmax function (used for converting scores into probabilities) and expectation-maximization algorithms used in statistical modeling. These functions inherently involve calculating probabilities that often require exponentiating large negative numbers, pushing the results toward the smallest representable float values.

Consider a large negative input value  $x$ , where we calculate  $e^x$ . If  $x$  is, say,  $-1000$ ,  $e^{-1000}$  is an extraordinarily small number, far smaller than the smallest positive number that a 64-bit float can represent without losing precision. When multiple such small terms are summed up in a numerator or denominator, they quickly underflow to zero. Even if the true ratio is a perfectly normal, manageable number (because the numerator and denominator are proportional), the direct computation in linear space fails because the intermediate values vanish.

When both the numerator ( $\sum e^{A_i}$ ) and the denominator ( $\sum e^{B_j}$ ) underflow to zero, the final operation becomes  $0/0$ . This operation is mathematically indeterminate, and the computational environment correctly registers this ambiguity by producing [NaN](#). This is a crucial failure point because it demonstrates a lack of [numerical stability](#). The problem is not with the mathematical model itself, but with the computer's ability to execute that model reliably using standard floating-point methods. Therefore, specialized techniques are required to transform the calculation into a domain where underflow does not occur.

## Reproducing the `double\_scalars` Instability

To provide a tangible demonstration of this numerical issue, we can construct a simple scenario using [NumPy](#) arrays. We define input values that, while mathematically sound, force the exponential function to produce results far outside the safe range of 64-bit precision. The following example involves large positive integers which, when multiplied by  $-3$  and exponentiated, drive the results toward numerical zero.

We aim to calculate the ratio of two sums of extremely small exponentials. When executed directly, the Python interpreter immediately raises the warning, indicating that the operation has failed to yield a determinate result. The code block below precisely replicates the conditions necessary to trigger the targeted [double\\_scalars](#) warning:

```
import numpy as np
```

```
#define two NumPy arrays containing large values
```

```
array1 = np.array(1)
```

```
array2 = np.array(1)
```

```
#perform complex mathematical operation: ratio of sums of exponentials
```

```
np.exp(-3*array1).sum() / np.exp(-3*array2).sum()
```

```
RuntimeWarning: invalid value encountered in double_scalars
```

Upon execution, the system issues the [RuntimeWarning](#). This occurs because the individual elements of the arrays (e.g.,  $e^{-3300}$  and  $e^{-3600}$ ) are calculated. These values are so minute that the finite precision of the 64-bit float representation cannot distinguish them from zero. Consequently, both `np.exp(-3*array1).sum()` and `np.exp(-3*array2).sum()` evaluate to 0.0, resulting in the mathematically invalid division  $0.0 / 0.0$ . This indeterminate operation is what Python flags specifically as an invalid value encountered in double-precision floating-point numbers.

## The Log-Sum-Exp Trick: The Robust Mathematical Solution

The definitive solution for handling ratios of exponentials, and thus resolving the `double\_scalars` warning in this context, is the application of the **Log-Sum-Exp (LSE) trick**. This is not a software patch but a fundamental mathematical transformation designed to move the calculation from the linear space (where underflow/overflow is common) into the logarithmic space, where numerical ranges are significantly compressed and manageable.

The LSE trick is based on a clever algebraic manipulation of the logarithm of a sum of

exponentials. Instead of directly calculating the potentially vanishing term  $\sum e^{x_i}$ , the method introduces a strategically chosen offset,  $a$  (typically the maximum value of  $x_i$ ), to scale the exponents. The core identity is expressed as:  $\ln(\sum e^{x_i}) = a + \ln(\sum e^{x_i - a})$ . By subtracting the maximum value  $a$  from all exponents, we ensure that the largest term inside the inner exponential is  $e^0 = 1$ . This rescaling prevents the exponents from becoming overwhelmingly negative, thus mitigating the risk of underflow to zero while simultaneously preventing overflow.

When applying the LSE trick to the specific problem of a ratio of sums of exponentials, we leverage the fundamental properties of logarithms. The original goal is to calculate  $R = \frac{\sum e^{A_i}}{\sum e^{B_j}}$ . By taking the logarithm of the ratio, the division is transformed into a subtraction:  $\ln(R) = \ln(\sum e^{A_i}) - \ln(\sum e^{B_j})$ . Since the terms on the right-hand side are now calculated using the stable Log-Sum-Exp form, we can safely compute the difference. Finally, we exponentiate the result to return to the original scale:  $R = e^{\ln(\sum e^{A_i}) - \ln(\sum e^{B_j})}$ . This transformation ensures high [numerical stability](#) throughout the entire process.

Fortunately, developers do not need to implement this complex trick manually. The SciPy library, designed specifically for scientific computing in Python, provides an optimized, pre-built function called [logsumexp](#) within its special functions module. This function handles the offsetting and logarithmic calculation efficiently, offering a clean, one-line solution to an otherwise challenging numerical problem.

## Implementing the Fix with SciPy's `logsumexp()`

To successfully fix the [RuntimeWarning](#) and perform the calculation robustly, we must import and utilize the [logsumexp](#) function from the `scipy.special` module. This function directly computes the logarithm of the sum of exponentials, guaranteeing that the intermediate calculations remain numerically stable, even when dealing with inputs spanning a vast dynamic range.

We modify the original code by replacing the direct exponentiation and summation with the [logsumexp](#) function, converting the division operation into a subtraction in log space. The final result is then exponentiated to obtain the value of the original ratio. Note that the input arrays remain unchanged; only the calculation methodology is altered to enhance [numerical stability](#):

```
import numpy as np
from scipy.special import logsumexp

#define two NumPy arrays (same as unstable example)
array1 = np.array()
array2 = np.array()
```

```
#perform complex mathematical operation using logsumexp for stability  
np.exp(logsumexp(-3*array1) - logsumexp(-3*array2))
```

```
2.7071782767869983e+195
```

The resulting output, `2.7071782767869983e+195`, is an extremely large number, demonstrating that the calculation was indeed mathematically valid and produced a massive value, which was simply unrepresentable using the direct method. Crucially, by employing the `logsumexp` approach, we have completely avoided the generation of [NaN](#) or the triggering of the [double\\_scalars](#) warning. This confirms that the instability was purely a computational artifact of the method used, not a flaw in the underlying mathematical model.

This technique is the industry standard for ensuring accuracy and reliability in scientific computing tasks involving dynamic ranges that push the limits of standard floating-point precision. Whenever you encounter the need to calculate the ratio of sums of exponentials (or the Softmax function), the use of specialized, numerically stable functions like `logsumexp()` is mandatory to preserve the integrity of your results.

## General Strategies for Robust Scientific Computing

While the Log-Sum-Exp trick is the targeted remedy for exponential ratios, professional scientific computing demands a broader approach to maintaining [numerical stability](#) across all operations. Adopting a set of best practices, particularly when utilizing powerful libraries like [NumPy](#) and [SciPy](#), minimizes the risk of encountering silent precision loss or debilitating runtime warnings.

Key strategies for ensuring robust computations include:

**Prioritizing Vectorization over Loops:** Always structure computations to utilize the highly optimized, vectorized operations provided by [NumPy](#). Vectorized functions are not only faster but are often implemented with internal safeguards against common numerical pitfalls, ensuring better handling of edge cases compared to standard Python loops.

**Selecting Appropriate Data Types:** While [double\\_scalars](#) (float64) offer excellent precision for most applications, awareness of your data's inherent range is vital. If input data requires massive numbers, ensure you normalize inputs to a manageable range (e.g., between 0 and 1) whenever possible before applying transcendental functions.

**Leveraging SciPy's Specialized Functions:** Before attempting to implement complex mathematical functions like Softmax, Sigmoid, or specialized statistical distributions manually, always check the SciPy special functions module. Functions such as `expit` (for sigmoid), `softmax`, and [logsumexp](#) are explicitly engineered by numerical analysts to incorporate the necessary scaling and log-space transformations internally, thus proactively preventing underflow and

overflow issues.

By consistently applying these principles, developers can transition from debugging hard-to-trace numerical errors to focusing on the scientific validity of their models. Specialized library functions are indispensable because they represent decades of collective experience in numerical analysis, guaranteeing that the mathematical foundation of your code is reliable, accurate, and resilient against the inherent limitations of floating-point arithmetic.

For those engaged in intensive computation, delving into the official documentation for SciPy is highly recommended. These resources provide detailed explanations of mathematical implementations and use cases for functions, equipping you with the tools needed to build truly robust computational models.