

Package: evaldrift (via r-universe)

October 22, 2026

Type Package

Title Paraphrase-Drift Scoring for LLM Eval Sets

Version 0.1.0

Maintainer James Verlander <scrapertweeter3@gmail.com>

Description Paraphrase-heavy eval sets drift over time: answers get rewritten, questions get reworded, and nobody re-runs the comparison that justified keeping both versions. eval_drift() scores two texts with three plain similarity measures (word-shingle Jaccard, character-trigram Jaccard, normalized Levenshtein similarity) and one blended drift score. No model calls, no network access, pure base R. Companion to the eval-methodology writing on PastAGI <<https://pastagi.com>>.

License MIT + file LICENSE

Encoding UTF-8

URL <https://pastagi.com>

BugReports <https://github.com/scrapertweeter3-prog/evaldrift/issues>

RoxygenNote 7.3.1

Repository <https://scrapertweeter3-prog.r-universe.dev>

Date/Publication 2026-10-21 23:19:50 UTC

RemoteUrl <https://github.com/scrapertweeter3-prog/evaldrift-r>

RemoteRef HEAD

RemoteSha 67b0f0530be28a47f11a57eba0680fd8f3549a7b

Contents

eval_drift	2
Index	3

eval_drift	<i>Paraphrase-drift score for a pair of texts</i>
------------	---

Description

Deterministic, dependency-free paraphrase-drift scoring so drifted eval pairs get flagged before a human ever looks at them.

Usage

```
eval_drift(a, b, weights = c(shingle = 0.40, trigram = 0.35, lev = 0.25), n = 3)
```

Arguments

a	First text (e.g. the original answer).
b	Second text (e.g. the rewritten answer).
weights	Blend weights for shingle, trigram, Levenshtein.
n	Word-shingle size.

Value

A list: shingle_jaccard, trigram_jaccard, levenshtein_sim, drift.

Examples

```
eval_drift("The model refuses the request.",  
           "The model declines to answer the request.")
```

Index

`eval_drift`, [2](#)