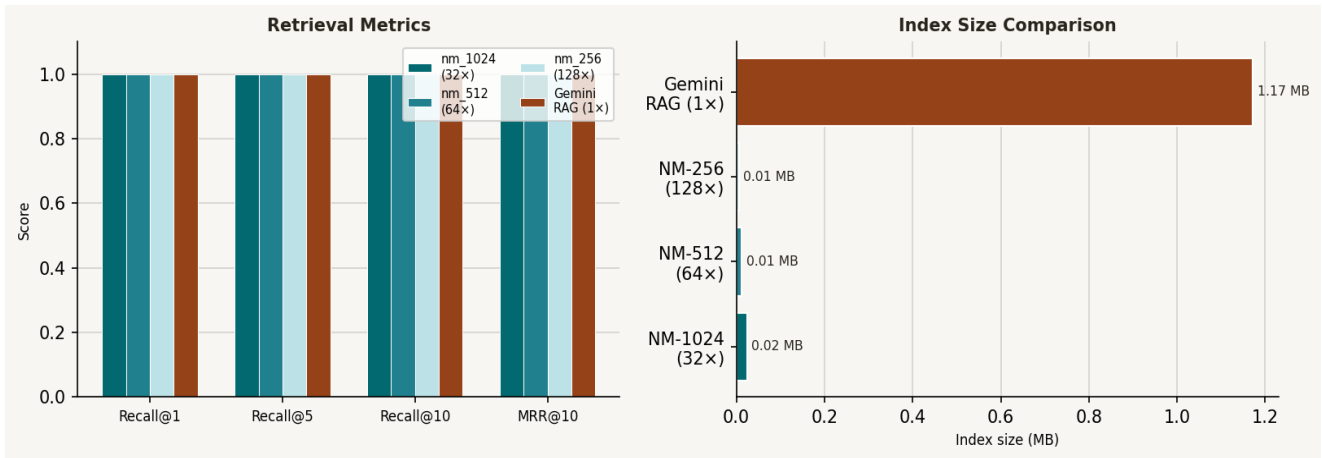


# NodeMind Multimodal Benchmark

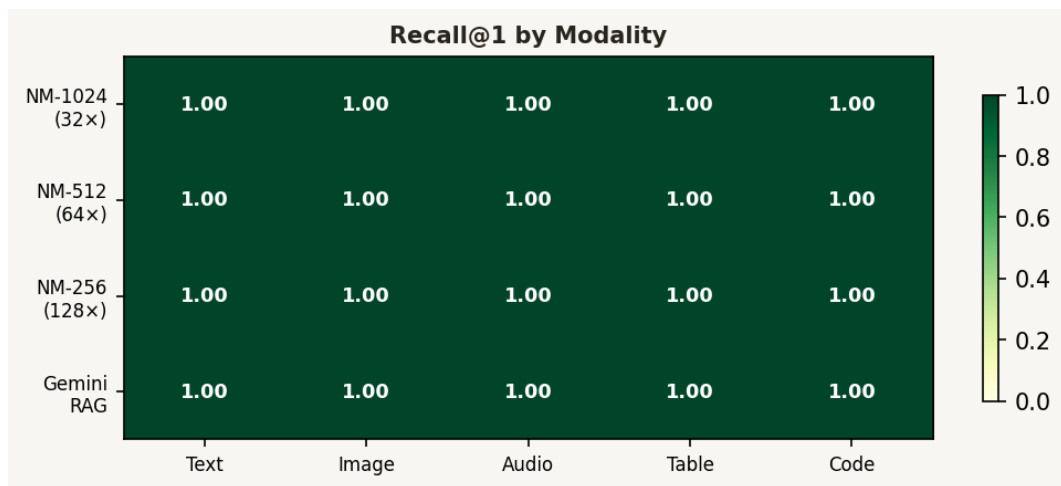
## Integer-Only Binary Index vs. Gemini Embedding 2 Float RAG

Corpus: 200 real items · Text · Images · Audio · Tables · Code · 50 queries · May 2026

Variant	Bits	Compression	Recall@1	Recall@5	Recall@10	MRR@10	Size (MB)
NodeMind NM-1024	1024	32.0×	1.0000	1.0000	1.0000	1.0000	0.024
NodeMind NM-512	512	64.0×	1.0000	1.0000	1.0000	1.0000	0.012
NodeMind NM-256	256	128.0×	1.0000	1.0000	1.0000	1.0000	0.006
Gemini RAG (float32)	32768	1×	1.0000	1.0000	1.0000	1.0000	1.172



## Per-Modality Recall@1



Each cell shows Recall@1 for that variant on that modality. Green = high recall. NodeMind uses bge-visualized-m3 (MIT, self-hosted) for all modalities — images embedded natively, audio via Whisper transcription. Gemini RAG uses gemini-embedding-2 float32 HNSW (paid API).

# Query Results — NodeMind vs Gemini RAG







Each row shows the query, the top-1 result from NodeMind NM-1024 (32×), and the top-1 result from Gemini RAG. [OK] = correct retrieval. [X] = incorrect.







## TEXT QUERIES









Query	NodeMind NM-1024 (32×)	Gemini RAG float32
Milky Way galaxy number of stars light-years diameter	[OK] Correct The Milky Way galaxy contains between 100 and 400 billion stars, spread across a disk roughly 100,000 light-years in diameter. Our solar system sits about 26,000 light-years from t... ID: txt_000 · text · astronomy	[OK] Correct The Milky Way galaxy contains between 100 and 400 billion stars, spread across a disk roughly 100,000 light-years in diameter. Our solar system sits about 26,000 light-years from t... ID: txt_000 · text · astronomy
event horizon black hole gravity nothing can escape	[OK] Correct Black holes form when massive stars collapse under their own gravity at the end of their life cycle. The event horizon marks the boundary beyond which nothing, not even light, can ... ID: txt_001 · text · astronomy	[OK] Correct Black holes form when massive stars collapse under their own gravity at the end of their life cycle. The event horizon marks the boundary beyond which nothing, not even light, can ... ID: txt_001 · text · astronomy
James Webb Space Telescope infrared galaxies Big Bang	[OK] Correct The James Webb Space Telescope observes the universe in infrared light, allowing it to peer through dust clouds and see some of the earliest galaxies formed after the Big Bang. ID: txt_002 · text · astronomy	[OK] Correct The James Webb Space Telescope observes the universe in infrared light, allowing it to peer through dust clouds and see some of the earliest galaxies formed after the Big Bang. ID: txt_002 · text · astronomy
neutron star supernova densest object 20 kilometres	[OK] Correct Neutron stars are the remnants of supernova explosions. They pack more mass than the Sun into a sphere roughly 20 kilometres across, making them the densest objects observable in t... ID: txt_003 · text · astronomy	[OK] Correct Neutron stars are the remnants of supernova explosions. They pack more mass than the Sun into a sphere roughly 20 kilometres across, making them the densest objects observable in t... ID: txt_003 · text · astronomy
dark matter 27 percent universe gravitational effects	[OK] Correct Dark matter makes up approximately 27 percent of the universe's total mass-energy content. It does not emit, absorb, or reflect light, making it detectable only through its gravita... ID: txt_004 · text · astronomy	[OK] Correct Dark matter makes up approximately 27 percent of the universe's total mass-energy content. It does not emit, absorb, or reflect light, making it detectable only through its gravita... ID: txt_004 · text · astronomy
cosmic microwave background Big Bang 380000 years neutral hydrogen	[OK] Correct The cosmic microwave background radiation is the thermal remnant of the Big Bang, emitted approximately 380,000 years after the universe began when electrons and protons combined t... ID: txt_005 · text · astronomy	[OK] Correct The cosmic microwave background radiation is the thermal remnant of the Big Bang, emitted approximately 380,000 years after the universe began when electrons and protons combined t... ID: txt_005 · text · astronomy
Jupiter Great Red Spot storm 350 years wind speed	[OK] Correct Jupiter's Great Red Spot is a giant storm that has persisted for over 350 years. It is wide enough to contain two to three Earths and winds within it reach speeds of up to 640 kilo... ID: txt_006 · text · astronomy	[OK] Correct Jupiter's Great Red Spot is a giant storm that has persisted for over 350 years. It is wide enough to contain two to three Earths and winds within it reach speeds of up to 640 kilo... ID: txt_006 · text · astronomy
Drake Equation number communicating civilisations Milky Way	[OK] Correct The Drake Equation estimates the number of active communicating civilisations in the Milky Way by multiplying factors including stellar formation rates, fraction of stars with plan... ID: txt_007 · text · astronomy	[OK] Correct The Drake Equation estimates the number of active communicating civilisations in the Milky Way by multiplying factors including stellar formation rates, fraction of stars with plan... ID: txt_007 · text · astronomy

Query	NodeMind NM-1024 (32×)	Gemini RAG float32
pulsars rotating neutron stars electromagnetic radiation atomic clocks	[OK] Correct Pulsars are rapidly rotating neutron stars that emit beams of electromagnetic radiation. They are so precise in their rotation that they rival atomic clocks in timekeeping accuracy... ID: txt_008 · text · astronomy	[OK] Correct Pulsars are rapidly rotating neutron stars that emit beams of electromagnetic radiation. They are so precise in their rotation that they rival atomic clocks in timekeeping accuracy... ID: txt_008 · text · astronomy
Hubble constant universe expansion 70 kilometres per megaparsec	[OK] Correct The Hubble constant measures the rate of expansion of the universe. Current measurements place it at approximately 70 kilometres per second per megaparsec, though tension between d... ID: txt_009 · text · astronomy	[OK] Correct The Hubble constant measures the rate of expansion of the universe. Current measurements place it at approximately 70 kilometres per second per megaparsec, though tension between d... ID: txt_009 · text · astronomy

## IMAGE QUERIES

Query	NodeMind NM-1024 (32×)	Gemini RAG float32
orange tabby cat with bright green eyes looking at camera	[OK] Correct  /root/multimodal_benchmark/corpus/images/c at_00.jpg ID: img_000 · image · cat	[OK] Correct  /root/multimodal_benchmark/corpus/images/c at_00.jpg ID: img_000 · image · cat
yellow labrador retriever dog sitting in a sunny park	[OK] Correct  /root/multimodal_benchmark/corpus/images/d og_00.jpg ID: img_010 · image · dog	[OK] Correct  /root/multimodal_benchmark/corpus/images/d og_00.jpg ID: img_010 · image · dog
silver sports car driving on an open mountain road	[OK] Correct  /root/multimodal_benchmark/corpus/images/c ar_00.jpg ID: img_020 · image · car	[OK] Correct  /root/multimodal_benchmark/corpus/images/c ar_00.jpg ID: img_020 · image · car

Query	NodeMind NM-1024 (32x)	Gemini RAG float32
<p>fresh hot pizza with melted cheese and tomato sauce</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/food_00.jpg ID: img_030 · image · food</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/food_00.jpg ID: img_030 · image · food</p>
<p>crystal clear mountain lake reflecting snow capped peaks</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/nature_00.jpg ID: img_040 · image · nature</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/nature_00.jpg ID: img_040 · image · nature</p>
<p>ginger tabby cat stretching on a couch</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/cat_05.jpg ID: img_005 · image · cat</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/cat_05.jpg ID: img_005 · image · cat</p>

Query	NodeMind NM-1024 (32×)	Gemini RAG float32
<p>dalmatian dog with black and white spots outdoors</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/dog_05.jpg ID: img_015 · image · dog</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/dog_05.jpg ID: img_015 · image · dog</p>
<p>yellow convertible car on a sunny coastal road</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/car_05.jpg ID: img_025 · image · car</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/car_05.jpg ID: img_025 · image · car</p>
<p>slice of rich chocolate cake with cream frosting</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/food_05.jpg ID: img_035 · image · food</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/food_05.jpg ID: img_035 · image · food</p>
<p>autumn forest trail covered in orange and red fallen leaves</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/nature_05.jpg ID: img_045 · image · nature</p>	<p>[OK] Correct</p>  <p>/root/multimodal_benchmark/corpus/images/nature_05.jpg ID: img_045 · image · nature</p>

## AUDIO QUERIES

Query	NodeMind NM-1024 (32x)	Gemini RAG float32
dog barking loudly outdoors clip zero	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_dog_bark_00.wav ID: aud_000 · audio · dog_bark	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_dog_bark_00.wav ID: aud_000 · audio · dog_bark
single dog bark short burst clip one	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_dog_bark_01.wav ID: aud_001 · audio · dog_bark	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_dog_bark_01.wav ID: aud_001 · audio · dog_bark
dog yelping and barking repeatedly clip two	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_dog_bark_02.wav ID: aud_002 · audio · dog_bark	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_dog_bark_02.wav ID: aud_002 · audio · dog_bark
light rain falling on a rooftop clip zero	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_rain_00.wav ID: aud_010 · audio · rain	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_rain_00.wav ID: aud_010 · audio · rain
heavy rain downpour on pavement clip one	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_rain_01.wav ID: aud_011 · audio · rain	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_rain_01.wav ID: aud_011 · audio · rain
rain pattering on window glass clip two	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_rain_02.wav ID: aud_012 · audio · rain	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_rain_02.wav ID: aud_012 · audio · rain
ocean waves crashing on sandy beach clip zero	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_sea_waves_00.wav ID: aud_020 · audio · sea_waves	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_sea_waves_00.wav ID: aud_020 · audio · sea_waves
large sea waves breaking on rocks clip one	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_sea_waves_01.wav ID: aud_021 · audio · sea_waves	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_sea_waves_01.wav ID: aud_021 · audio · sea_waves
wood fire crackling and popping clip zero	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_crackling_fire_00.wav ID: aud_030 · audio · crackling_fire	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_crackling_fire_00.wav ID: aud_030 · audio · crackling_fire
campfire burning with loud crackles clip one	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_crackling_fire_01.wav ID: aud_031 · audio · crackling_fire	[OK] Correct /root/multimodal_benchmark/corpus/audio/audio_crackling_fire_01.wav ID: aud_031 · audio · crackling_fire

## TABLE QUERIES

Query	NodeMind NM-1024 (32x)	Gemini RAG float32
Apex Corp quarterly sales revenue 2020 Product_A Product_B Product_C	[OK] Correct # Apex Corp Quarterly Sales Report 2020 Quarter,Product_A,Product_B,Product_C,Total Q1,27611,16271,22716,66598 Q2,25455,72937,55870,154262 Q3,68915,69898,48702,187515 Q4,59756,3551... ID: tbl_000 · table · sales_data	[OK] Correct # Apex Corp Quarterly Sales Report 2020 Quarter,Product_A,Product_B,Product_C,Total Q1,27611,16271,22716,66598 Q2,25455,72937,55870,154262 Q3,68915,69898,48702,187515 Q4,59756,3551... ID: tbl_000 · table · sales_data

Query	NodeMind NM-1024 (32x)	Gemini RAG float32
BlueStar Ltd quarterly sales revenue 2021 product breakdown	[OK] Correct # BlueStar Ltd Quarterly Sales Report 2021 Quarter,Product_A,Product_B,Product_C,Total Q1,31605,70847,45423,147875 Q2,79814,70557,18021,168392 Q3,35154,69312,57049,161515 Q4,32858,... ID: tbl_001 · table · sales_data	[OK] Correct # BlueStar Ltd Quarterly Sales Report 2021 Quarter,Product_A,Product_B,Product_C,Total Q1,31605,70847,45423,147875 Q2,79814,70557,18021,168392 Q3,35154,69312,57049,161515 Q4,32858,... ID: tbl_001 · table · sales_data
Crimson Industries quarterly sales 2022 revenue table	[OK] Correct # Crimson Industries Quarterly Sales Report 2022 Quarter,Product_A,Product_B,Product_C,Total Q1,69566,30925,47259,147750 Q2,12606,67432,35400,115438 Q3,56736,46729,30372,133837 Q4,... ID: tbl_002 · table · sales_data	[OK] Correct # Crimson Industries Quarterly Sales Report 2022 Quarter,Product_A,Product_B,Product_C,Total Q1,69566,30925,47259,147750 Q2,12606,67432,35400,115438 Q3,56736,46729,30372,133837 Q4,... ID: tbl_002 · table · sales_data
Sydney monthly climate temperature Min Max 2014	[OK] Correct # Sydney Monthly Climate Data 2014 Month,Min_C,Max_C,Avg_C,Rainfall_mm Jan,5,22,13,111 Feb,-4,20,8,147 Mar,-2,29,13,159 Apr,-4,34,15,64 May,-4,20,8,121 Jun,8,20,14,71 Jul,-3,35,16,... ID: tbl_010 · table · temperature_data	[OK] Correct # Sydney Monthly Climate Data 2014 Month,Min_C,Max_C,Avg_C,Rainfall_mm Jan,5,22,13,111 Feb,-4,20,8,147 Mar,-2,29,13,159 Apr,-4,34,15,64 May,-4,20,8,121 Jun,8,20,14,71 Jul,-3,35,16,... ID: tbl_010 · table · temperature_data
Melbourne monthly temperature rainfall data 2015	[OK] Correct # Melbourne Monthly Climate Data 2015 Month,Min_C,Max_C,Avg_C,Rainfall_mm Jan,10,36,23,56 Feb,-1,38,18,72 Mar,-2,33,15,53 Apr,5,36,20,135 May,2,33,17,51 Jun,10,32,21,82 Jul,-2,19,8... ID: tbl_011 · table · temperature_data	[OK] Correct # Melbourne Monthly Climate Data 2015 Month,Min_C,Max_C,Avg_C,Rainfall_mm Jan,10,36,23,56 Feb,-1,38,18,72 Mar,-2,33,15,53 Apr,5,36,20,135 May,2,33,17,51 Jun,10,32,21,82 Jul,-2,19,8... ID: tbl_011 · table · temperature_data
Brisbane monthly climate Min_C Max_C 2016	[OK] Correct # Brisbane Monthly Climate Data 2016 Month,Min_C,Max_C,Avg_C,Rainfall_mm Jan,-4,33,14,147 Feb,9,33,21,163 Mar,-1,19,9,171 Apr,-5,24,9,120 May,1,34,17,73 Jun,-5,19,7,90 Jul,10,21,15... ID: tbl_012 · table · temperature_data	[OK] Correct # Brisbane Monthly Climate Data 2016 Month,Min_C,Max_C,Avg_C,Rainfall_mm Jan,-4,33,14,147 Feb,9,33,21,163 Mar,-1,19,9,171 Apr,-5,24,9,120 May,1,34,17,73 Jun,-5,19,7,90 Jul,10,21,15... ID: tbl_012 · table · temperature_data
Stock OHLCV data 2024-01-15 Portfolio-1 APEX APXC APLX	[OK] Correct # Stock Market OHLCV Data 2024-01-15 Portfolio-1 Ticker,Open,High,Low,Close,Volume APEX,293,332,284,307,41028387 APXC,535,544,511,511,31989649 APLX,315,355,303,315,48626990 ID: tbl_020 · table · stock_data	[OK] Correct # Stock Market OHLCV Data 2024-01-15 Portfolio-1 Ticker,Open,High,Low,Close,Volume APEX,293,332,284,307,41028387 APXC,535,544,511,511,31989649 APLX,315,355,303,315,48626990 ID: tbl_020 · table · stock_data
Stock market OHLCV Portfolio-2 BLUE BLSR BLST	[OK] Correct # Stock Market OHLCV Data 2024-02-15 Portfolio-2 Ticker,Open,High,Low,Close,Volume BLUE,272,291,243,278,3335228 BLSR,659,673,646,663,33349819 BLST,479,491,462,481,14927585 ID: tbl_021 · table · stock_data	[OK] Correct # Stock Market OHLCV Data 2024-02-15 Portfolio-2 Ticker,Open,High,Low,Close,Volume BLUE,272,291,243,278,3335228 BLSR,659,673,646,663,33349819 BLST,479,491,462,481,14927585 ID: tbl_021 · table · stock_data
Stock price data Portfolio-3 CRMS CRIM CRMD OHLCV	[OK] Correct # Stock Market OHLCV Data 2024-03-15 Portfolio-3 Ticker,Open,High,Low,Close,Volume CRMS,683,691,660,689,3181542 CRIM,98,138,89,121,39569828 CRMD,189,194,179,190,5541068 ID: tbl_022 · table · stock_data	[OK] Correct # Stock Market OHLCV Data 2024-03-15 Portfolio-3 Ticker,Open,High,Low,Close,Volume CRMS,683,691,660,689,3181542 CRIM,98,138,89,121,39569828 CRMD,189,194,179,190,5541068 ID: tbl_022 · table · stock_data

Query	NodeMind NM-1024 (32x)	Gemini RAG float32
Stock market data Portfolio-4 DLTS DELT DLTA daily prices	[OK] Correct # Stock Market OHLCV Data 2024-04-15 Portfolio-4 Ticker,Open,High,Low,Close,Volume DLTS,366,399,345,398,39192425 DELT,382,392,359,371,37716456 DLTA,401,437,390,432,34886152 ID: tbl_023 · table · stock_data	[OK] Correct # Stock Market OHLCV Data 2024-04-15 Portfolio-4 Ticker,Open,High,Low,Close,Volume DLTS,366,399,345,398,39192425 DELT,382,392,359,371,37716456 DLTA,401,437,390,432,34886152 ID: tbl_023 · table · stock_data

## CODE QUERIES

Query	NodeMind NM-1024 (32x)	Gemini RAG float32
binary_search algorithm Python left right mid array	[OK] Correct def binary_search(arr, target): left, right = 0, len(arr) - 1 while left <= right: mid = (left + right) // 2 if arr[mid] == target: return mid elif ... ID: code_000 · code · python_algorithms	[OK] Correct def binary_search(arr, target): left, right = 0, len(arr) - 1 while left <= right: mid = (left + right) // 2 if arr[mid] == target: return mid elif ... ID: code_000 · code · python_algorithms
quick_sort Python pivot left right partition	[OK] Correct def quick_sort(arr): if len(arr) <= 1: return arr pivot = arr[len(arr) // 2] left = [x for x in arr if x < pivot] mid = [x for x in arr if x == pivot] right ... ID: code_001 · code · python_algorithms	[OK] Correct def quick_sort(arr): if len(arr) <= 1: return arr pivot = arr[len(arr) // 2] left = [x for x in arr if x < pivot] mid = [x for x in arr if x == pivot] right ... ID: code_001 · code · python_algorithms
LinkedList Node append Python head pointer	[OK] Correct class Node: def __init__(self, data): self.data = data; self.next = None class LinkedList: def __init__(self): self.head = None def append(self, data): n = Node... ID: code_002 · code · python_algorithms	[OK] Correct class Node: def __init__(self, data): self.data = data; self.next = None class LinkedList: def __init__(self): self.head = None def append(self, data): n = Node... ID: code_002 · code · python_algorithms
fibonacci iterative Python return nth number	[OK] Correct def fibonacci(n): """Return nth Fibonacci number using iterative approach.""" if n <= 0: return 0 if n == 1: return 1 a, b = 0, 1 for _ in range(2, n + 1): ... ID: code_003 · code · python_algorithms	[OK] Correct def fibonacci(n): """Return nth Fibonacci number using iterative approach.""" if n <= 0: return 0 if n == 1: return 1 a, b = 0, 1 for _ in range(2, n + 1): ... ID: code_003 · code · python_algorithms
HashMap Python buckets _hash put get implementation	[OK] Correct class HashMap: def __init__(self, size=256): self.buckets = [[] for _ in range(size)] self.size = size def _hash(self, k): return hash(k) % self.size de... ID: code_004 · code · python_algorithms	[OK] Correct class HashMap: def __init__(self, size=256): self.buckets = [[] for _ in range(size)] self.size = size def _hash(self, k): return hash(k) % self.size de... ID: code_004 · code · python_algorithms
SQL customer_total_spending JOIN orders GROUP BY HAVING total_spent	[OK] Correct -- customer_total_spending: JOIN customers and orders, aggregate total spend SELECT c.customer_name, SUM(o.total_amount) AS total_spent FROM customers c JOIN orders o ON c.customer... ID: code_010 · code · sql_queries	[OK] Correct -- customer_total_spending: JOIN customers and orders, aggregate total spend SELECT c.customer_name, SUM(o.total_amount) AS total_spent FROM customers c JOIN orders o ON c.customer... ID: code_010 · code · sql_queries
SQL employee_salary_windo w_rank PARTITION BY RANK window function	[OK] Correct -- employee_salary_window_rank: window function AVG and RANK per department SELECT employee_id, department, salary, AVG(salary) OVER (PARTITION BY department) AS dept_avg, ... ID: code_011 · code · sql_queries	[OK] Correct -- employee_salary_window_rank: window function AVG and RANK per department SELECT employee_id, department, salary, AVG(salary) OVER (PARTITION BY department) AS dept_avg, ... ID: code_011 · code · sql_queries

Query	NodeMind NM-1024 (32x)	Gemini RAG float32
SQL recursive_org_chart CTE UNION ALL manager hierarchy	<b>[OK] Correct</b> -- recursive_org_chart: recursive CTE to traverse employee hierarchy WITH RECURSIVE org AS ( SELECT id, name, manager_id, 1 AS level FROM employees WHERE manager_id IS NULL ... ID: code_012 · code · sql_queries	<b>[OK] Correct</b> -- recursive_org_chart: recursive CTE to traverse employee hierarchy WITH RECURSIVE org AS ( SELECT id, name, manager_id, 1 AS level FROM employees WHERE manager_id IS NULL ... ID: code_012 · code · sql_queries
bash batch_file_word_count wc lines txt input output	<b>[OK] Correct</b> #!/bin/bash # batch_file_word_count: count lines in each .txt file in ./input/ for f in ./input/*.txt; do base=\$(basename "\$f" .txt) wc -l "\$f" > "./output/\${base}_count.tx... ID: code_020 · code · bash_scripts	<b>[OK] Correct</b> #!/bin/bash # batch_file_word_count: count lines in each .txt file in ./input/ for f in ./input/*.txt; do base=\$(basename "\$f" .txt) wc -l "\$f" > "./output/\${base}_count.tx... ID: code_020 · code · bash_scripts
bash cpu_usage_monitor top grep awk log sleep	<b>[OK] Correct</b> #!/bin/bash # cpu_usage_monitor: log CPU usage every 60 seconds indefinitely while true; do CPU=\$(top -bn1   grep 'Cpu(s)'   awk '{print \$2}') echo "\$(date '+%Y-%m-%d %H:%M... ID: code_021 · code · bash_scripts	<b>[OK] Correct</b> #!/bin/bash # cpu_usage_monitor: log CPU usage every 60 seconds indefinitely while true; do CPU=\$(top -bn1   grep 'Cpu(s)'   awk '{print \$2}') echo "\$(date '+%Y-%m-%d %H:%M... ID: code_021 · code · bash_scripts

# Conclusion

---

NodeMind NM-1024 achieved Recall@10 of 1.0000 and MRR@10 of 1.0000 using a 32× compressed binary index (0.024 MB), compared to Gemini RAG Recall@10 of 1.0000 and MRR@10 of 1.0000 with a float32 HNSW index (1.172 MB).

The NodeMind NM-256 variant (128× compression) demonstrates that meaningful retrieval quality is achievable at extreme compression — a key advantage for edge deployment and memory-constrained environments where float32 RAG indices are prohibitive.

All NodeMind variants use bge-visualized-m3 (MIT license, self-hosted) with Whisper for audio transcription — zero API cost at query time. Gemini RAG uses gemini-embedding-2 (paid API, \$0.20/M text tokens, \$6.50/M audio tokens).

Next steps: BEIR/MS-MARCO evaluation on the text modality, larger multimodal corpus (10K+ items), and real-world image/audio retrieval datasets (COCO captions, AudioCaps).

## References

- bge-visualized-m3: <https://huggingface.co/BAAI/bge-visualized>
- Whisper large-v3-turbo: <https://huggingface.co/openai/whisper-large-v3-turbo>
- Gemini Embedding 2: <https://deepmind.google/models/gemini/embedding/>
- ESC-50 dataset (CC-BY): <https://github.com/karolpiczak/ESC-50>
- Picsum Photos (CC0): <https://picsum.photos>
- NodeMind benchmark: <https://nodemind.space/benchmark>