DPRG RBNV Chat Record - March 25, 2025

00:04:55.983,00:04:58.983
Paul Bouchier: Steve described his work using RAG AI to understand the DPRG mail list archives that are on the web site

00:06:32.020,00:06:35.020
Paul Bouchier: Steve indexed the archives in both a dense & sparse index, which each have their own strengths

00:19:59.981,00:20:02.981
Paul Bouchier: Steve demonstrated different types of search: dense, sparse. He explained that a RAG search would include results from the internet in addition to what it found in the DPRG email archives

00:21:01.287,00:21:04.287
Paul Bouchier: The search of indexes has no chance of hallucinations because it's only looking at the email archive

00:24:32.228,00:24:35.228
Paul Bouchier: The temperature setting affects creativity in assessing relevance of returned results - higher means it will be more liberal about what it returns.

00:26:07.592,00:26:10.592
Paul Bouchier: The reset command resets the context, so it gets rid of filters active in the previous context

00:32:53.756,00:32:56.756
Paul Bouchier: You can ask it to summarize results

00:33:45.843,00:33:48.843
Paul Bouchier: Steve's working on a web interface to the search engine.

00:34:02.175,00:34:05.175
Carl Ott: Hey this has to be worth some kind of star!

00:34:06.533,00:34:09.533
Steve Edwards: https://github.com/sedwardstx/DprgArchiveAgent

00:34:51.592,00:34:54.592
Paul Bouchier: Steve posted the URL (above) of his code for the DPRG archive agent

00:35:32.200,00:35:35.200
Jim F - CalgaryAB: @Carl how about Silver Star (Gold Stars for Robotics)

00:36:57.773,00:37:00.773
Paul Bouchier: The search engine doesn't require any special compute resources.

00:47:15.546,00:47:18.546
ed mart: Steve ... what would you add in the future ?

00:49:38.497,00:49:41.497
Paul Bouchier: Steve did a one-time ingest. At present there's no ongoing ingestion

00:50:04.865,00:50:07.865
Paul Bouchier: Steve mooted the idea of auto-ingesting into Discord

00:54:18.635,00:54:21.635
Paul Bouchier: If DPRG offered this AI search as a feature, and purchased tokens to support it, the expectation is that the cost would be in the few dollars a month for light usage.

00:55:01.143,00:55:04.143
Paul Bouchier: Karim: OpenAI offers discounts to non-profits - would it be interesting to DPRG?

00:58:27.795,00:58:30.795
Paul Bouchier: Another way of using it is to find a relevant post then go to the DPRG archive to read it.

01:01:26.828,01:01:29.828
Steve Edwards: to answer questions about sparse vs dense indexes, here are the pro and con for each.

01:01:30.818,01:01:33.818
Steve Edwards: Dense Indexes
Pros:

High Accuracy and Recall: Dense indexes include every vector, which ensures that the nearest neighbor search is exhaustive and yields high-quality, precise results.

Consistency: Since every vector is indexed, the search results are consistent across queries, with minimal chance of missing a relevant vector.

Simplicity: The implementation can be straightforward because there’s no need to decide which vectors to omit or approximate.

01:01:38.219,01:01:41.219
Steve Edwards: Cons:

Memory Intensive: Storing an index entry for every vector can be very memory-consuming, especially for large-scale datasets.

Slower for Massive Datasets: As the dataset grows, the time required to scan the entire index can increase, leading to slower query performance.

Scalability Challenges: High resource usage makes it harder to scale the database efficiently when working with millions or billions of vectors.

Sparse Indexes
Pros:

01:01:50.328,01:01:53.328
Steve Edwards: Reduced Memory Footprint: By indexing only a subset or by using approximations, sparse indexes consume significantly less memory.

Faster Query Performance: With fewer index entries to scan, queries can be executed more rapidly, which is advantageous in time-sensitive applications.

Better Scalability: Lower resource demands allow sparse indexes to scale more effectively for very large datasets.

01:01:58.432,01:02:01.432
Steve Edwards: Cons:

Potential Loss in Recall: Since not every vector is directly indexed, some relevant vectors might be omitted during the search, which can reduce the overall accuracy.

Complexity in Design: Deciding on which parts of the data to index (or how to effectively approximate the search) introduces additional complexity in system design and tuning.

Trade-off Between Speed and Accuracy: While faster, the search might sacrifice precision and relevance of results compared to a dense approach.

01:03:23.499,01:03:26.499
Paul Bouchier: GOLD STAR for Tom for showing line following running

01:04:18.216,01:04:21.216
Paul Bouchier: Pat showed his robot identifying the closest can in a live demo.

01:05:37.428,01:05:40.428
Paul Bouchier: Pat is using opencv running on the RPi to identify cans from a stereo camera.

01:06:05.368,01:06:08.368

Paul Bouchier: Pat is getting results at 30fps

01:10:38.682,01:10:41.682
Paul Bouchier: Pat is splitting the image into LAB (Luminance & A & B blended color channels

01:13:49.131,01:13:52.131
Paul Bouchier: I think Steve deserves a black star for major technical advances

01:20:23.659,01:20:26.659
Paul Bouchier: Pat will use a Xiao RP2040 to read encoders & send odometry to the RPi

01:23:44.383,01:23:47.383
Paul Bouchier: BLACK STAR for Steve for his email search AI work

01:25:37.115,01:25:40.115
Paul Bouchier: Mike has been going to a ROS2 navigator class, & made a 6-can world, in gazebo

01:29:11.356,01:29:14.356
Ted Meyers: got to go

01:30:21.197,01:30:24.197
Paul Bouchier: Mike showed a robot localized in the 6-can world using its lidar & amcl in ROS nav2

01:34:00.106,01:34:03.106
Paul Bouchier: Mike showed placing cans in gazebo & showed the robot driving around the course while avoiding cans

01:36:07.672,01:36:10.672
Paul Bouchier: The robot pushed one of the cans around on its way somewhere.

01:51:03.960,01:51:06.960
Paul Bouchier: GOLD STAR for Mike for showing a robot moving in simulation, navigating among cans

01:52:41.846,01:52:44.846
Paul Bouchier: Mike showed his work on converting nodes to ROS2 lifecycle control

01:54:49.526,01:54:52.526
Paul Bouchier: He starts & stops timers when nodes transition between active & inactive lifecycle states

02:01:23.848,02:01:26.848

Paul Bouchier: Mike showed his downstairs robot in an arena, with the spinning lidar showing the arena wall and the left & right ToF sensors displaying cans they see in rviz. The left & right sensors are publishing PointCloud messages, and the center ToF sensors is pubishing a Range message

02:02:56.695,02:02:59.695
Steve Edwards: great meeting, i need to drop so looking forward to the recording to catch the rest. thanks all!

02:05:52.896,02:05:55.896
Pat Caron: See you next week guys