



Zoom Meeting ID
503 616 9422

Seven AI Patterns

`Vienna is calling. Answer the damn phone`

Matt Navarre · navarre.ai · free classes every Thursday

About Me



- **Passionate trainer, Expert developer**
- **Founded 3 companies**
- **FBA Awards, conference sessions, meetups, YouTube vids**
- **ClarisTalk AI podcast with Cris Ippolite**
 - **20 years of FileMaker Podcasts**

My One Request

Make A free appointment with me to talk
about AI in your solutions

www.navarre.ai

● Agenda

1. Hello World – Invoice Summary
2. AI Models – Which, when and Why?
3. fmSearchResults AI
 - a. Basic Semantic Search
 - b. ‘Enriched’ Search
4. Error Finder – on Load
5. Reporting – Can AI do math?
6. Image Description (and PDF and URL)
7. Building & iterating WebViewer layouts

● ‘Perform Semantic Find’

Perform Semantic Find + Insert Embedding ship out of the box.



Great at

- Paragraphs of prose
- Product descriptions, notes, resumes
- Support tickets, meeting transcripts
- “What does this mean?” lookups



Quietly bad at

- Names (“Bob Smith”)
- Company names and addresses
- SKUs, short codes, tight labels
- Dates, prices, numeric ranges

Most data in a FileMaker solution lives up here ^ ^

● Semantically clever. Numerically blind.

THE HONEST LIMIT



Meaning, similarity, paraphrase

“law office audit” matches “Financial Audit Preparation Law Office” because the WORDS mean similar things.



Dates, amounts, ranges, exact strings

“Projects in Q2” vs a due-date of 2026-05-14 — a vector will not know those are the same thing. Numbers are not meanings.



Right tool rule:

semantic search for meaning + a plain WHERE clause for dates and amounts.

Stop trying to make one tool do both. That's the #1 mistake I see with new FileMaker AI work.

● Default Config

QUERY

“Robert Smyth”



DATABASE HAS

Bob Smith

● Enrich & Embed the data



SYSTEM PROMPT · person

"Generate realistic name variants for use in a semantic search index. Always include: the original name, LastName-FirstName form, accent-free spelling, common nicknames or diminutives, transliteration variants if relevant. Return a plain text list – one variant per line."

Everything goes into the embedding. The record side is now fat with signal.

● Enrich & Embed the Query :: Set to Set Search

ENRICHED RECORD

Bob Smith
Robert Smith
Smith, Robert
Bobby / Rob / Robb
Smyth / Smithe
ABC Technology
bob@abctech.com



ENRICHED QUERY

Robert Smyth
Rob Smyth
R. Smyth
Smyth, Robert
robertt@abc-tech.com
237 More Park Dr, CA
Alphabet Tek?

Orthogonal variants meet in the middle. Typos, nicknames and transliterations fall on the same point in vector space.

● Expand to companies, projects, products, etc.



ENTITY → COMPANIES

WHAT WE EMBED

- Original org name
- Abbreviations & acronyms
- With / without suffix
- List of staff names
- Project description
- Invoice Description
- Product info

REAL-WORLD MATCH

User types:

PA Family Law

Finds:

Pennsylvania Divorce Resources

User types:

ACME secondary office update project

Finds:

A C M E Satellite Lighting & Ergonomics update

● Teach the index what “cheap” means



ENTITY → PRODUCTS

Users search “cheap red shoes.” The database stores \$39.99 / crimson / Nike.

PRICE TIERS

Shoes:

< \$50 → cheap, budget
\$50–150 → mid-range
> \$200 → luxury, premium

Cars:

< \$20k → basic, economy
> \$50k → luxury, flagship

COLOR FAMILIES

Red →

crimson, cherry,
cardinal, burgundy,
wine, garnet, brick,
tomato, rose, maroon

Silver → platinum,
ash, pewter, titanium

CATEGORY SYNONYMS

Shoes →

sneakers, trainers,
kicks, running shoes

SUV →

crossover, 4x4,
sport utility

“Cheap” is not a price. It's a category-relative judgment. Bake that in at write-time.

● Re-enrich and re-embed



- Gen + Embed cost: basically free

~100

tokens in

the search string + a compact system
prompt

~200

tokens out

5–10 name or query variants

small

Dumb embed is fine

gpt-4-mini, haiku, local models.
No frontier needed.