

A wide-angle photograph of a glacial lake. In the foreground, several large, blue-tinted icebergs of various sizes are scattered across the dark water. Behind them, a range of mountains is covered in thick snow and ice, with some rocky peaks visible through the white snow. The sky above is overcast with soft, grey clouds.

Building Relevance Formulas with LLMs

-Kristian Aune, Vespa.ai

2025-04-23, Haystack



KRISTIAN AUNE: CO-FOUNDER VESPA.AI AND HEAD OF CUSTOMER SUCCESS

20 years working with application owners, developers, PMs, Ops - search, recommendation and personalization systems.

Spun out of Yahoo, October 2023 - \$31M funding.

Open source 2017, Apache 2.0 license. Available on Vespa Cloud since 2021.



Qwant⁺

yahoo!

RavenPack

Elicit

Spotify R&D

OTTO

Vinted

okcupid



What is important?

S's criteria:

- Safe
- Reliable
- Prestige (brand)

Other hard criteria:

- Pre-owned (to make it affordable)
- Automatic transmission (so I can drive it)
- Minimum engine size (M is picky)



Search 🔍

Heart Save Search | Clear Filters Manufacturer Certified AWD/4WD (8) Rear Wheel Drive Front Wheel Drive Uninsured (45)

Location (1)

Distance ZIP Code

Nationwide ZIP Code

Include Extended Delivery (1)

Year (1)

Minimum Year Maximum Year

Make, Model & Trim (1)

Vehicle 1 Any Make

New
 Used
 Manufacturer Certified (1)
 Third-Party Certified (1)

Drive Type (1)

Price & Payments (1)

Min Price Max Price

Max Payment Personalize Payment

Ex \$600/mo.

HUGE AVAILABILITY OF METRIS VANS! Show payments on listings

Heart Save Search | Clear Filters Manufacturer Certified AWD/4WD (8) Rear Wheel Drive Front Wheel Drive Uninsured (45)

My Wallet Sort By: Relevance

172,350 Results



Certified 2021 Jeep Gladiator Sport

26,048 miles

Jeep Certified

\$32,848 GOOD PRICE

See estimated payment

Sponsored by Coon Rapids Chrysler Dodge Jeep Ram
(952) 213-0796

[Check Availability](#)

Certified 2022 Volvo XC90 T8 Inscription Expression

31,443 miles

Volvo Certified Plug-In Hybrid

\$53,295 GOOD PRICE

See estimated payment

Sponsored by Rusnak Volvo Cars
(626) 667-9377

[Check Availability](#)

Certified 2024 Porsche Macan

3,622 miles

Porsche Certified

\$61,995 GREAT PRICE

Sponsored by Serra Auto Campus
(517) 940-7835

Online Paperwork

[Check Availability](#)

Picked for You



Certified 2023 Mercedes-Benz Metris Passenger

8,977 miles

\$42,499

GREAT PRICE

HUGE AVAILABILITY OF METRIS VANS!

Certified 2022 Kia Sportage LX

21,383 miles

\$22,830

GOOD PRICE

Certified 2023 Nissan Rogue SV

8,513 miles

\$23,798

GREAT PRICE

Certified 2023 Nissan Titan PRO-4X w/ Pro-4x Convenience Package

13,680 miles

\$45,995

GOOD PRICE

Certified 2022 Ford F250 Super Duty

8,529 miles

\$76,000

GREAT PRICE

Show payments on listings

Helpful Links[Best Car Deals](#)[Best Truck Deals](#)[Best SUV Deals](#)[Buy Cars Online](#)**Car Research & News**[Research & Information](#)[New & Future Models](#)[Trim Comparisons](#)**Car Buying Tools**[Buying and Selling Tips](#)[Car Payment Calculator](#)[Affordability Calculator](#)[Car Lease Calculator](#)

See if you can save on
car insurance on
[KBB.com](#)



SPONSORED

Pre-qualify for a car loan
with no credit score
impact



Get a Kelley Blue Book
instant cash offer.

Used Cars for Sale Near You

[Browse by Make](#) | [Shop Used Cars](#) | [Advanced Search](#)

Any Make

Any Model

ZIP Code

Search

The most rugged Honda
we've ever dreamt up.



2026 Honda Passport

Learn More

Sponsored by HONDA

Used Cars by Price

\$450

per month

Continue

Add your down payment and trade-in to see cars in your budget.



Buying a Used Car

Sure, used cars may not come with that nice “new-car smell,” but a used car value is almost always the better financial decision. On average, drivers save about \$200 monthly on used car payments versus new car payments. Not to mention how quickly new cars depreciate once they leave the lot! That fancy new model will drop 35% of its value in the first three years of its life, while you can still find a used car with tip-top technology, high-end packages, and aftermarket additions you don’t have to buy yourself. Whether you’re buying for yourself or a new driver in the family, a used car can be a great deal when shopping at dealerships or [private sellers](#).

[Shop Used Cars](#)

which car brands are considered safe

Pro Search Sources 10

2 tasks ↗

 Yahoo
10 Safest Car Brands in America
- Yahoo Finance

 Wheel Front Ltd. Liability...
Study Reveals Safest Car Brands
Over the Last Decade | Wheel...

 Park+
Safest Car Brands: As per 2025
Ranking - Park+

 USN 
+7 sources

Several car brands are widely recognized for their commitment to safety, based on crash test ratings, advanced safety features, and real-world performance. Here are some of the top brands considered safe:

Top Safe Car Brands

1. Volvo

- Known for pioneering safety innovations, every Volvo model released in the past decade has achieved a 5-star NCAP rating [2](#) [4](#).
- Features include blind-spot monitoring, forward collision warning with pedestrian detection, and Pilot Assist [3](#).

2. Subaru

- Subaru models like the Outback and Impreza consistently earn high safety marks, supported by its EyeSight driver-assist technology [3](#).
- All Subaru models released in the last decade have achieved a perfect 5-star NCAP rating [2](#).

3. Tesla

- Tesla vehicles such as the Model X and Model Y are equipped with advanced driver assistance systems (ADAS) and have maintained perfect 5-star NCAP ratings [1](#) [2](#).

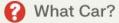
which car brands are reliable

Pro Search Sources 15

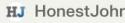
2 tasks ↗



Ranked: The Most Reliable Car Brands in 2025 - Visual...



Britain's most reliable cars 2024 | plus the ones that'll let you...



Most reliable car brands 2025 | Honest John



+12 sources

The most reliable car brands are consistently highlighted across various surveys and rankings for their durability, engineering quality, and customer satisfaction. Here are the top brands based on recent data:

Top Reliable Car Brands in 2025

1. Lexus

- Lexus frequently ranks as the most reliable brand due to its robust engineering and low incidence of issues. Models like the NX and RX are particularly dependable [1](#) [3](#) [5](#).

2. Toyota

- Known for its legendary reliability, Toyota offers durable vehicles across all price ranges, from the Aygo X to the RAV4. It consistently ranks near the top in reliability studies [1](#) [3](#) [5](#).

3. Subaru

- Subaru leads in reliability due to shared components across its models, reducing risks of new issues. The Forester and Impreza are standout performers [4](#) [6](#).

4. Mazda

- Mazda combines reliability with sporty design. Models like the CX-5 and MX-5 are highly rated for dependability [1](#) [4](#).

show great brands for the medium to high income families

Pro Search Sources 20

2 tasks ↗

Dourado Luxury Car - D...
List of Luxury Car Brands Best
for Families

CR Consumer Reports
Best Cars of the Year: 10 Top
Picks of 2025 - Consumer...

Yahoo
Top 8 Car Brands Driven By the
Wealthy - Yahoo Finance

+17 sources

For medium- to high-income families, several car brands stand out for their balance of luxury, reliability, safety, and practicality. These brands cater to families seeking comfort, advanced features, and premium experiences without compromising on functionality.

Great Car Brands for Medium- to High-Income Families

Luxury and Premium Options

1. Volvo

- Known for safety and Scandinavian design, models like the XC90 and XC60 offer spacious interiors, advanced safety features (e.g., collision avoidance), and luxurious comfort [1](#) [5](#).

2. Mercedes-Benz

- Models like the GLS and GLE provide a premium driving experience with high-end materials, advanced driver-assistance systems, and spacious interiors suited for families [1](#) [5](#).

3. BMW

- Family-friendly SUVs like the X5 and X7 combine luxury with performance. They feature cutting-edge technology, spacious cabins, and dynamic handling [1](#) [5](#).

4. Audi

- Vehicles such as the Q7 and Q8 offer precision engineering, ample space, Quattro all-wheel drive for stability, and family-oriented luxury [1](#) [5](#).

5. Lexus

- Renowned for reliability and comfort, models like the RX 350 and GX are equipped with

after some research ...

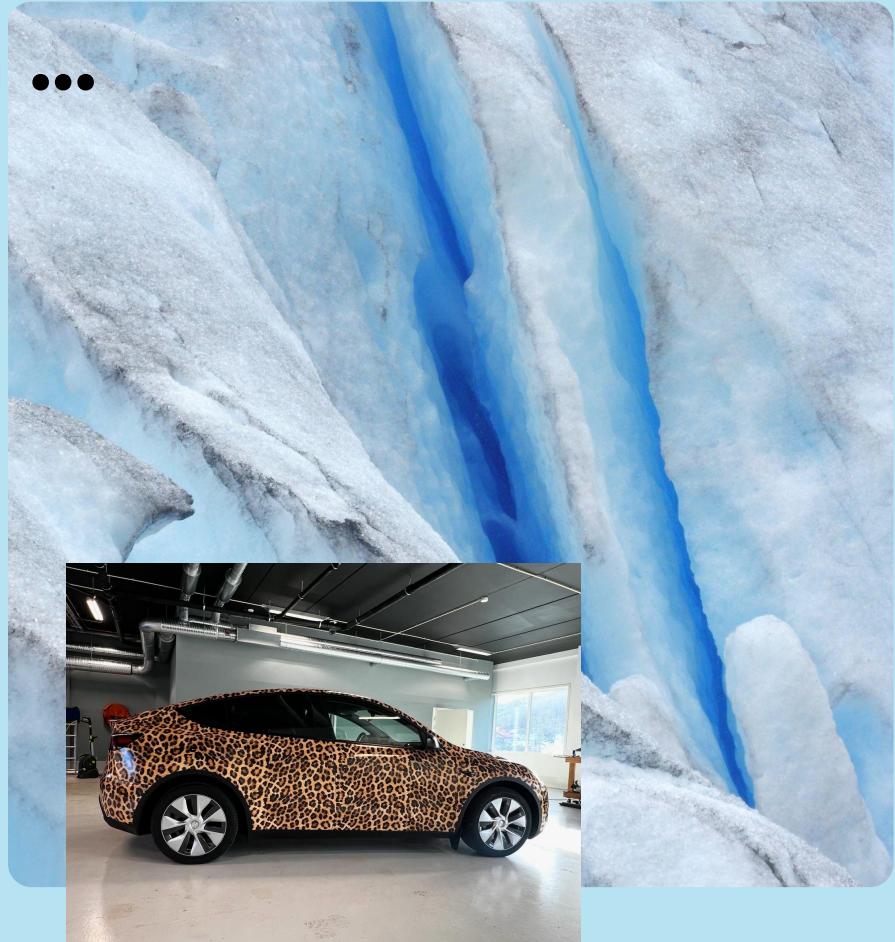
Safe → **Volvo, Subaru, Tesla**

Reliable → **Lexus, Toyota, Subaru, Mazda**

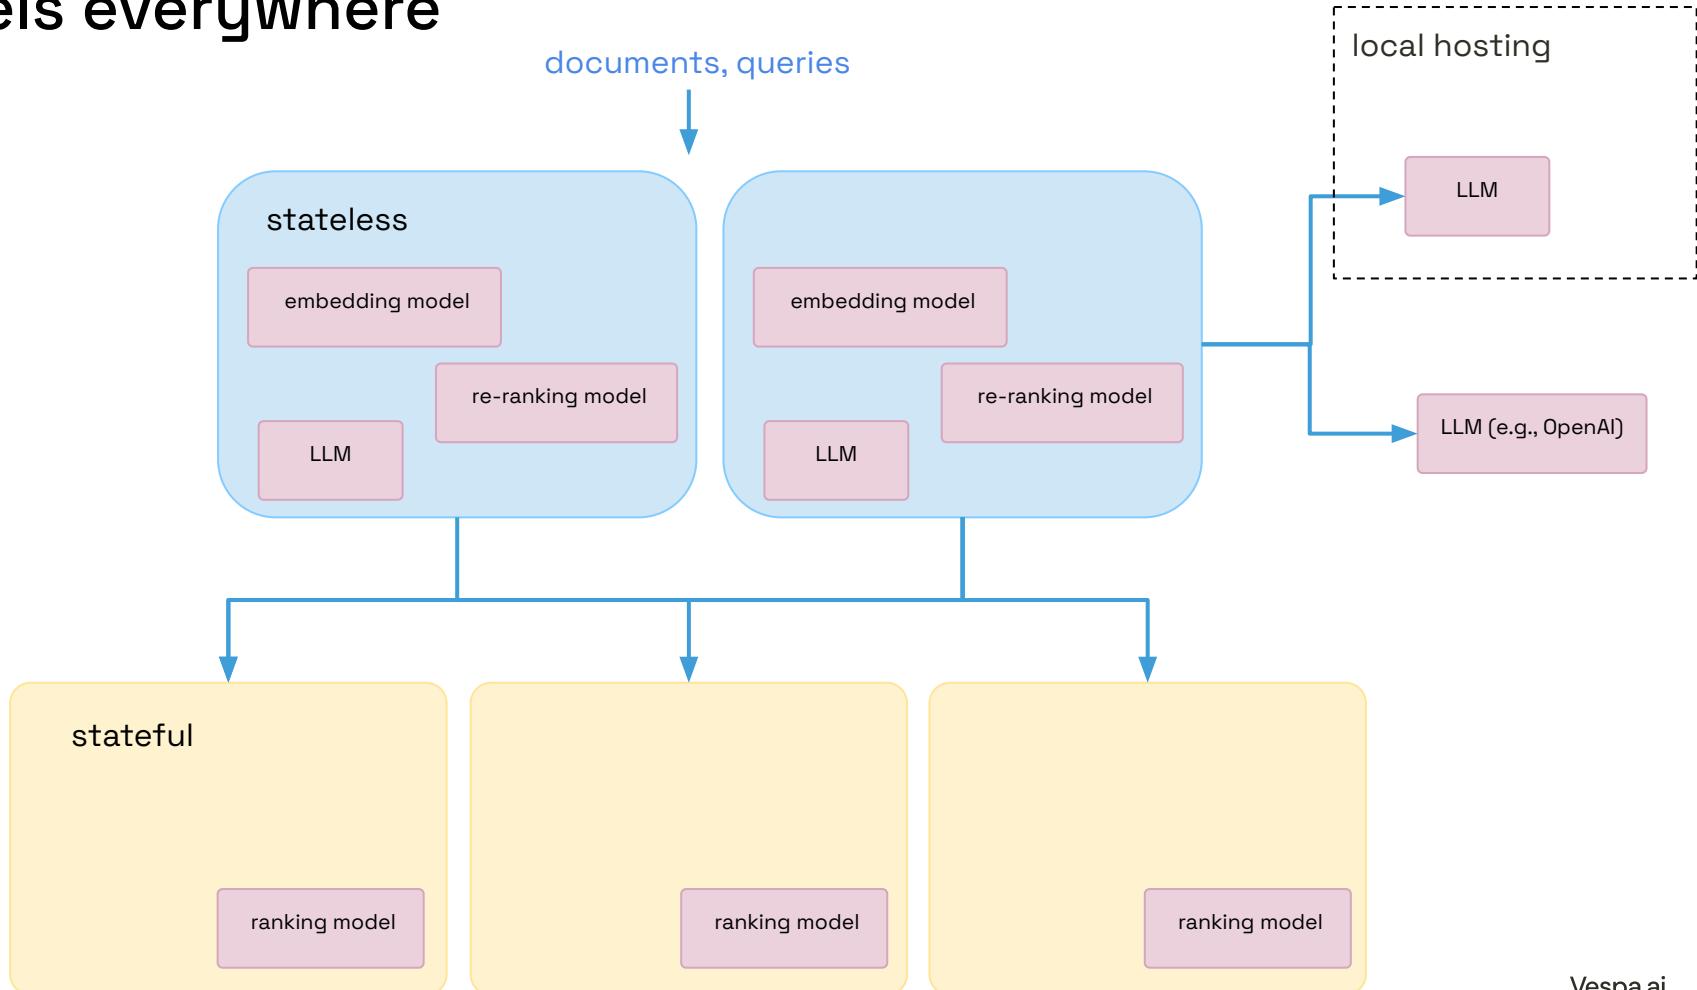
Prestige → **Volvo, Mercedes, BMW, Audi, Lexus**

How users express their preferences and criteria
is very open

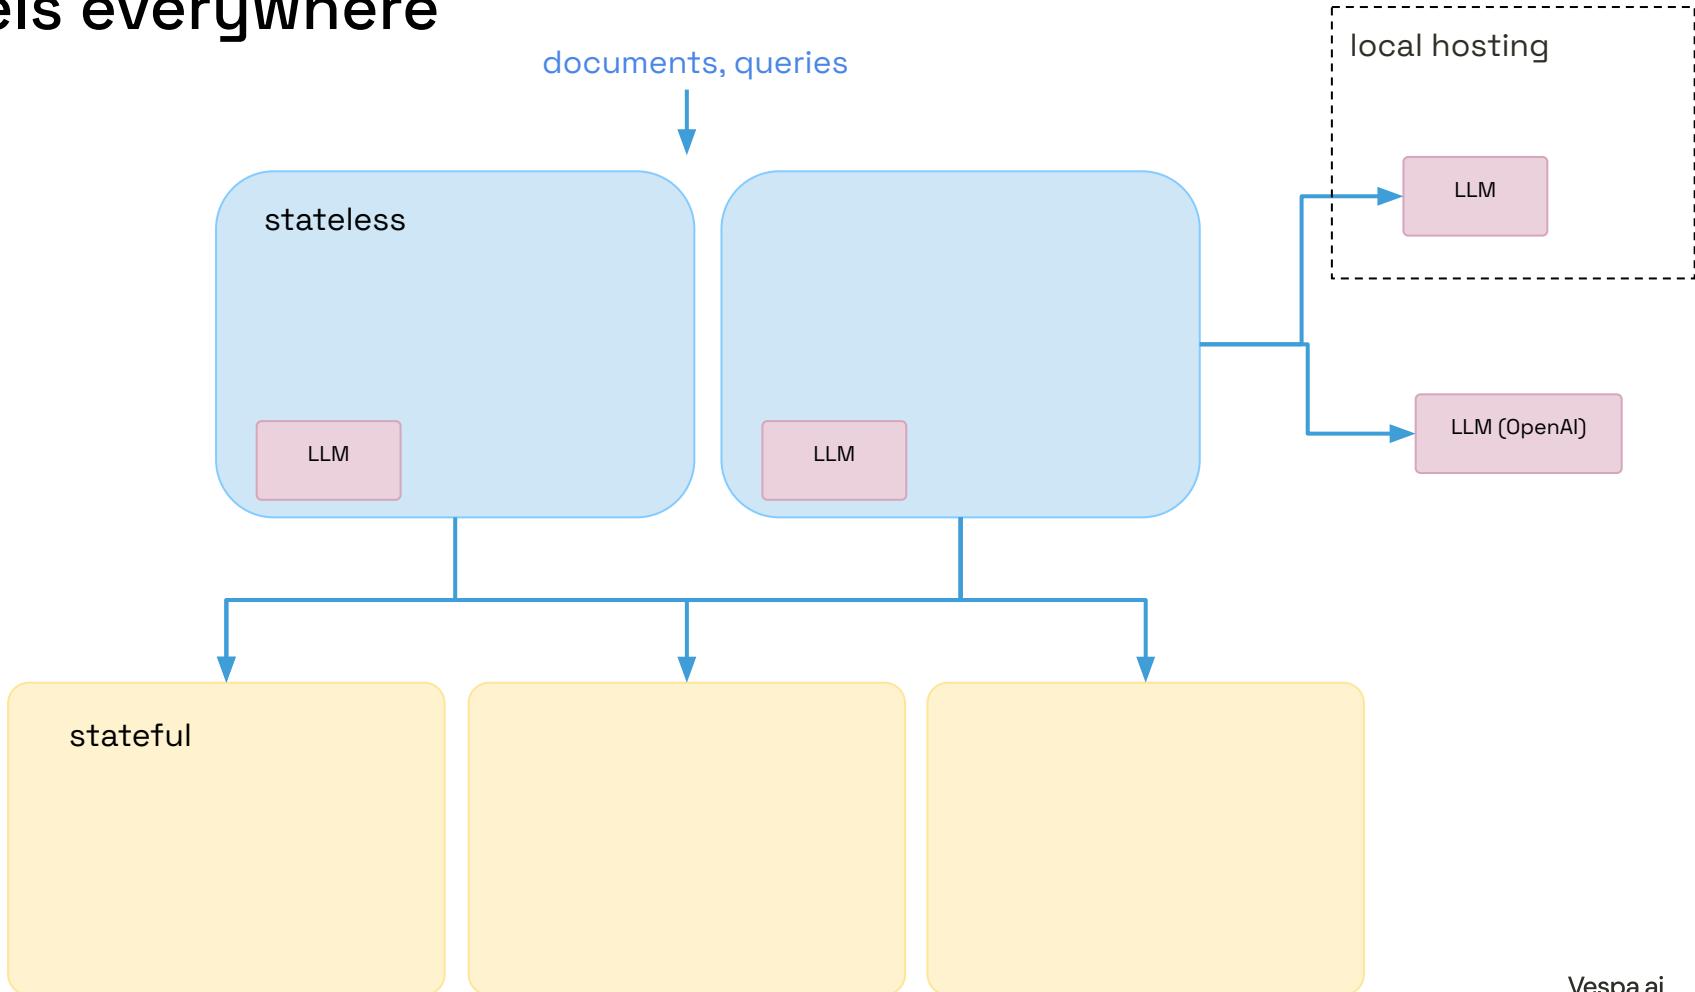
- therefore, hard to model into the schema
- perfect problem for an LLM!



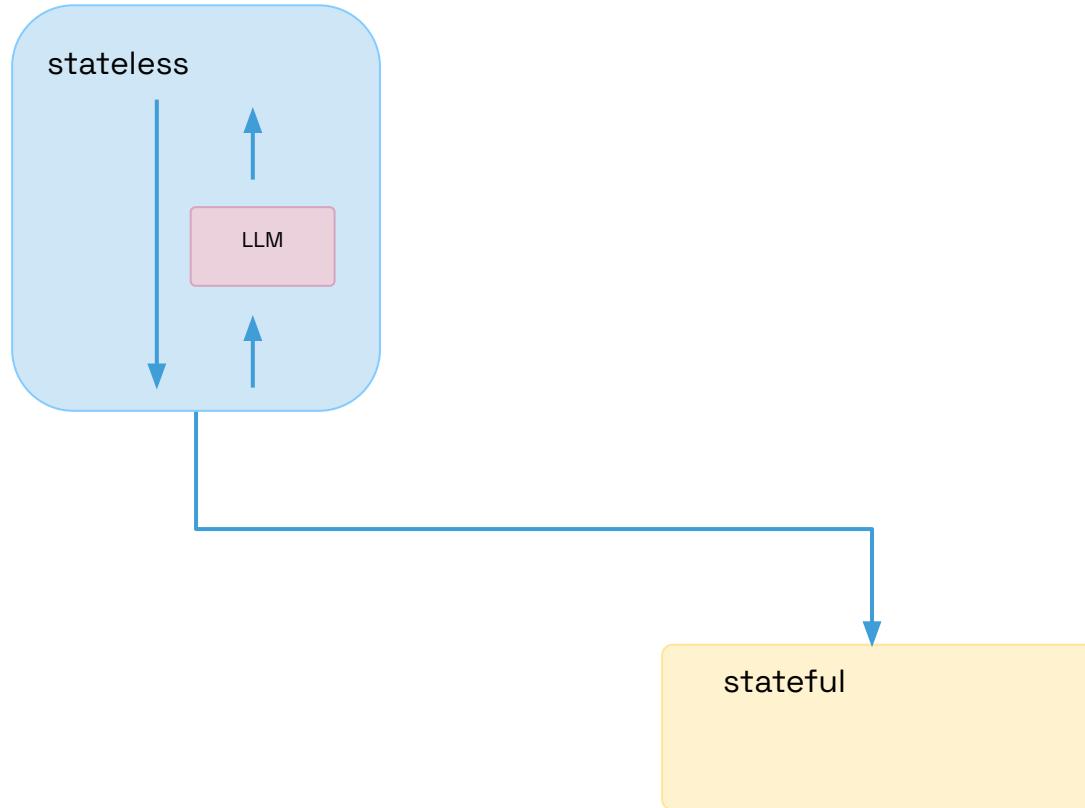
Models everywhere



Models everywhere



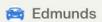
RAG (result summarization)



find the best used volvo close to charlottesville with a \$10000 budget

Perplexity Sources 10

3 tasks ↗



Used Volvo for Sale in Charlottesville, VA - Edmunds



Used cars for under \$10,000 in Charlottesville, VA - CarGurus



Used Cars Under \$10000 for Sale in Charlottesville, VA - ...



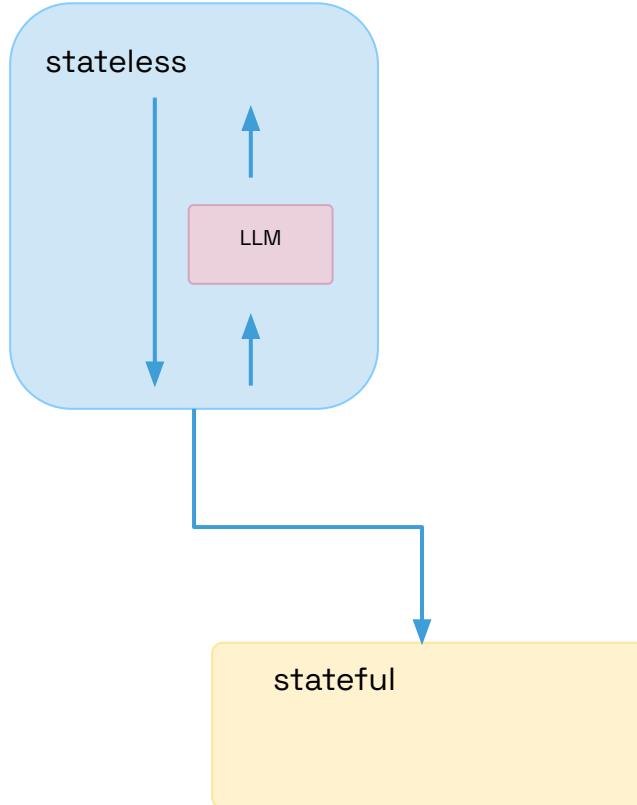
+7 sources

Best Used Volvos Near Charlottesville Under \$10,000

Several used Volvo models are available within a \$10,000 budget near Charlottesville, VA. The best options balance price, mileage, accident history, and overall value. Here are the top choices based on current listings:

Model & Year	Price	Mileage	Accident History	Location (Distance)	Notable Features
2015 Volvo XC60 T5 Drive-E Premier	\$7,999	Not listed	Not listed	Charlottesville, VA (local)	SUV, modern safety features
2014 Volvo XC60 T6 4dr SUV	\$9,374	95,486	1 accident, 1 owner	Chantilly, VA (86 mi)	6-cyl, clean title, personal use
2013 Volvo XC60 3.2 4dr SUV	\$7,495	101,766	2 accidents, 3 owners	Stafford, VA (69 mi)	6-cyl, clean title, personal use
2013 Volvo XC90 3.2 4dr SUV	\$5,995	153,708	No accidents, 3 owners	Arlington, VA (100 mi)	6-cyl, clean title, personal use
2012 Volvo C70 T5 Convertible	\$6,250	95,805	No accidents, 2 owners	Woodbridge, VA (83 mi)	Convertible, clean title
2012 Volvo C70 T5 Convertible	\$7,500	103,000	No accidents, 1 owner	Exclusive Imports (90 mi)	Convertible, clean title

Re-ranking results



Cars for Sale in Charlottesville, VA 22901

Buckle up! Car prices may increase, but we're here to help you navigate affordable payments.

Learn more about tariffs

Search

Location

Distance ZIP Code

50 Miles 22901

Include Extended Delivery

Year

Minimum Year Maximum Year

Any Any

Make, Model & Trim

Vehicle 1

Any Make

Condition

New (1,814) Used (1,285) Manufacturer Certified (169) Third-Party Certified (15)

Drive Type

Price & Payments

Min Price Max Price
\$4,500 \$181,600

[Save Search](#) | [Clear Filters](#) | [Purchase Online \(22\)](#) | [AWD/](#)

[Recent Price Drop](#)  [MALLOY](#) 

 **New 2025 Lincoln Nautilus Reserve w/ Jet Apps** **\$66,455** MSRP

CMA's Colonial Lincoln  (434) 424-0852 • [Check Availability](#)

3,099 Results

 **Used 2022 Kia Sportage EX** **\$18,721**  [See estimated payment](#)

66,731 miles  **\$100,570** MSRP

Sponsored by Steven Nissan 30.28 mi. away • (540) 217-4756 [Online Paperwork](#)

 **New 2025 GMC Hummer EV 2X** **\$100,570** MSRP

8 miles  **\$100,570** MSRP

Sponsored by CMA'S Colonial Buick GMC 1.86 mi. away • (434) 424-0852 [Online Paperwork](#)

 **Used 2020 Jeep Cherokee Latitude** **95,829 miles**

 [Contact Dealer For Price](#) [See estimated payment](#)

Sponsored by Battlefield Chevrolet 36.26 mi. away • (540) 445-0528

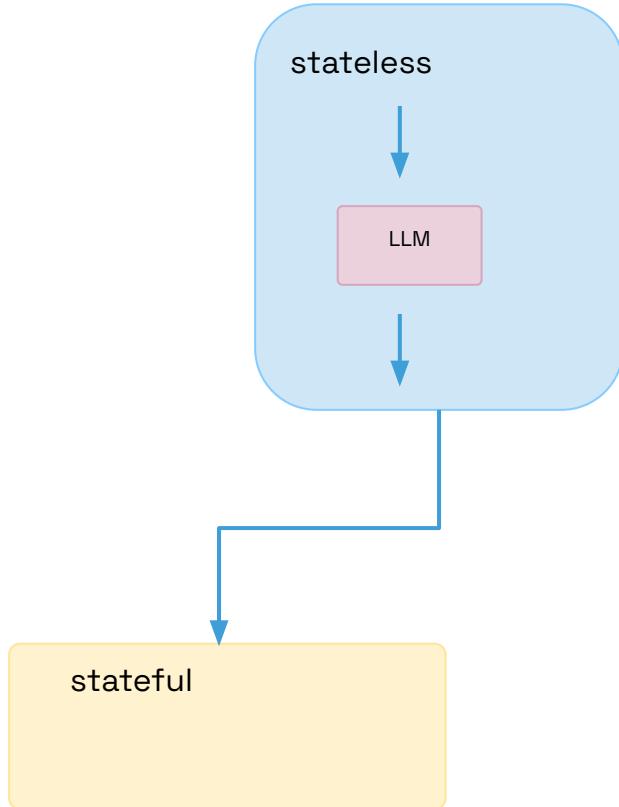
 





Indexing documents



Cars for Sale in Charlottesville, VA 22901

Buckle up! Car prices may increase, but we're here to help you navigate affordable payments.

Search 🔍

Location

Distance: 50 Miles ZIP Code: 22901

Include Extended Delivery ⓘ

Year

Minimum Year: Any Maximum Year: Any

Make, Model & Trim

Vehicle 1: Any Make

Condition

New (1,814) Used (1,285) Manufacturer Certified (169) ⓘ Third-Party Certified (15) ⓘ

Drive Type

Price & Payments

Min Price: \$4,500 Max Price: \$181,600

Save Search | Clear Filters 'automatic sport utility' × Purchase Online (2)

Sponsored by CMA's Colonial Lincoln

New 2025 Lincoln
\$66,455 MSRP

High Demand

CMA's Colonial Lincoln (434) 424-0852 • Check Availability

3,099 Results

Recent Price Drop

Used 2022 Kia Sportage EX
66,731 miles
\$18,721 GREAT PRICE
See estimated payment

New 2025 G

8 miles Electric
\$100,570

Sponsored by Steven Nissan

30.28 mi. away • (540) 217-4756

Online Paperwork

Check Availability ... **Check Availability** ... **Check Availability** ...

Raw Data



Enrichment Tasks

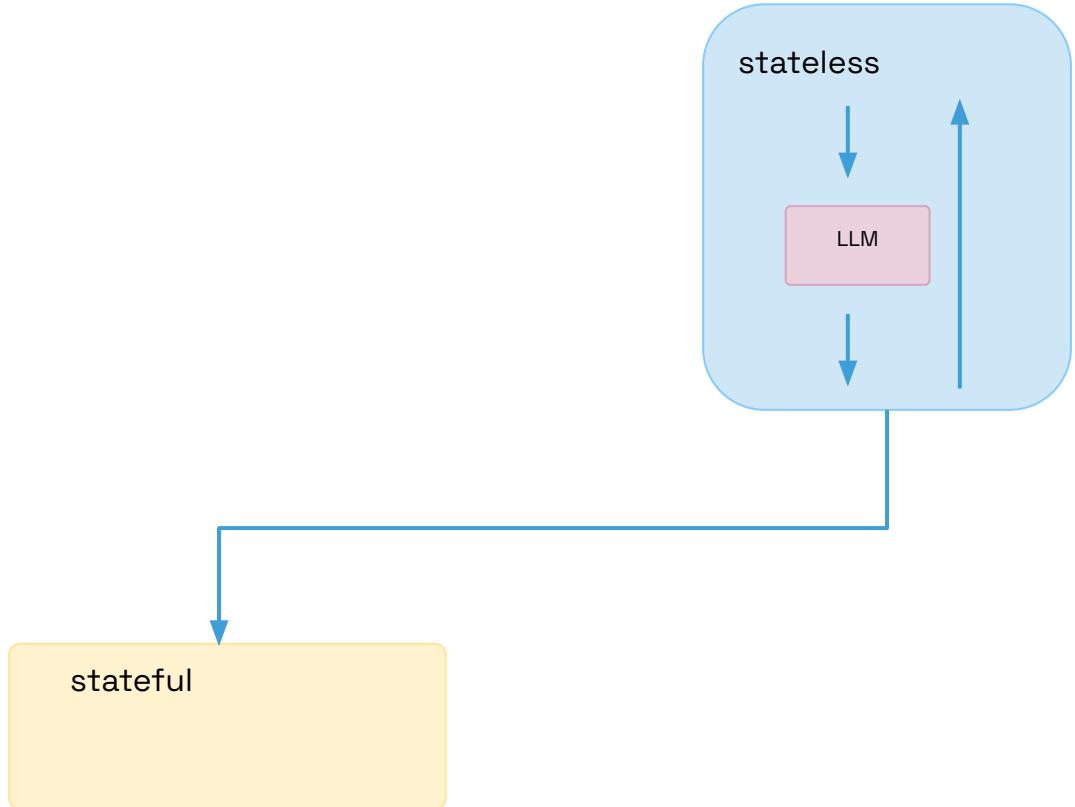


Improved Search Experience



```
{  
  "id": "id:product:product::6ce19adf-6badfcd40df2",  
  "fields": {  
    "mpg": 74.3,  
    "mileage": 56022,  
    "price": 8450,  
    "make": "hongqi",  
    "model": "e-hs9",  
    "year": 2016,  
    "make_alt": ["honchi", "honci", "hon-qi"],  
    "queries": [  
      "I want a safe car for my family",  
      "I prefer reliable cars with a strong US brand"  
    ]  
  }  
}
```

Query relevance



Cars for Sale in Charlottesville, VA 22901

Buckle up! Car prices may increase, but we're here to help you navigate affordable payments.

[Learn more about tariffs](#)

[Save Search](#) | [Clear Filters](#)

'automatic sport utility'

Purchase Online (22)

AWD/4WD (2,773)

Rear Wheel Drive

Location

Distance ZIP Code

50 Miles 22901

Include Extended Delivery

Year

Minimum Year Maximum Year

Any Any

Make, Model & Trim

Vehicle 1

Any Make

Condition

- New (1,814)
- Used (1,285)
- Manufacturer Certified (169)
- Third-Party Certified (15)

Drive Type

Price & Payments

Min Price Max Price

\$4,500 \$181,600

Sponsored by CMA's Colonial Lincoln



New 2025 Lincoln Nautilus Reserve w/ Jet Appearance
\$66,455 MSRP

CMA's Colonial Lincoln
(434) 424-0852 • [Check Availability](#)

3,099 Results



Used 2022 Kia Sportage EX

66,731 miles

\$18,721 GREAT PRICE

[See estimated payment](#)

Sponsored by Steven Nissan
30.28 mi. away • (540) 217-4756

[Check Availability](#)



New 2025 GMC Hummer EV 2 X

8 miles

Electric

\$100,570 MSRP

Sponsored by CMA'S Colonial Buick GMC
1.86 mi. away • (434) 424-0852
[Online Paperwork](#)

[Check Availability](#)



Used 2020 Jeep Cherokee Latitude

95,829 miles

[Contact Dealer For Price](#)

[See estimated payment](#)

Sponsored by Battlefield Chevrolet
36.26 mi. away • (540) 445-0528

[Check Availability](#)



My Wallet

✓ Relevance

- Price - Lowest
- Price - Highest
- Distance - Closest
- Date - Oldest
- Date - Newest
- Mileage - Lowest
- Mileage - Highest
- Year - Oldest
- Year - Newest

i prefer audi over bmw

Do you have a preference for a specific transmission or fuel type?

automatic transmission

Do you have a preference for a specific fuel type?

i like safe, german cars

Would you prefer a specific fuel type or transmission for the car?

no ford please

Do you have a preference for a specific transmission or fuel type?

i prefer audi over bmw

Do you have a preference for a specific transmission or fuel type?

automatic transmission

Do you have a preference for a specific fuel type?

I'd like to decrease the importance of bmw from 1.5 to 1.2.

Do you have a preference for a specific fuel type?

Tell me about your car preferences...



Facets

Make

Ford	17,965
Vw	15,157
Opel	13,632
Mercedes	13,119
Bmw	10,781
Audi	10,668
Toyota	6,738
Skoda	6,267
Hyundai	4,860

Transmission

Manual	56,445
Semi-Auto	22,677
Automatic	20,056
Other	9

Search Results

Your Preferences: (click to adjust)

Age +2.0 Audi +1.6 Automatic +2.0 Bmw +1.6 Diesel -1.0 Ford +2.0
Mercedes +1.5 Mileage +2.0 Saab +1.5 Volkswagen +1.5 Volvo +1.5



audi a1

PRICE:

\$21,400

YEAR:

2020

MILEAGE:

7 miles

TRANSMISSION:

automatic

FUEL TYPE:

petrol



audi q2

PRICE:

\$26,500

YEAR:

2020

MILEAGE:

10 miles

TRANSMISSION:

automatic

FUEL TYPE:

petrol

All

- Arts & Photography
- Biographies & Memoirs
- Business & Money
- Calendars
- Children's Books
- Comics & Graphic Novels
- Computers & Technology
- Cookbooks, Food & Wine

Crafts, Hobbies &

Rufus beta

show me good books on internet search

Here are some books that discuss different aspects of internet search:

The Net Delusion: The Dark Side of Internet...
by Evgeny Morozov

4.3 ★★★★★ (139)
Hardcover
\$22⁹¹ (\$0.98/oz)

List: \$27.95
FREE delivery Fri, May 9 to Norway on \$49 of eligible items

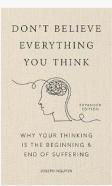
Examines the dark side of internet freedom, written by Evgeny Morozov.



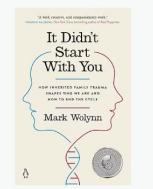
Ask Rufus a question



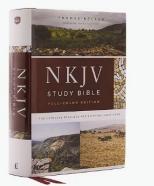
40% off Limited time deal
Baby University Complete "For Babies" Board Book Set



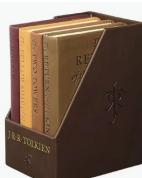
42% off Limited time deal
Don't Believe Everything You Think: Expanded



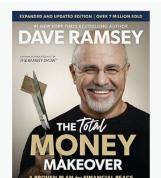
24% off Limited time deal
It Didn't Start with You: How Inherited Family



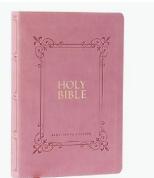
56% off Limited time deal
NKJV Study Bible, HardCover, Burgundy, Full-



50% off Limited time deal
The Hobbit And The Lord Of The Rings: Deluxe Pocket



62% off Limited time deal
The Total Money Makeover Updated and Expanded: A



64% off Limited time deal
KJV Holy Bible: Large Print with 53,000 Center-Column



The Net Delusion: The Dark Side of Internet...
by Evgeny Morozov

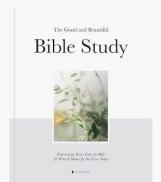
4.3 ★★★★★ (139)
Hardcover
\$22⁹¹ (\$0.98/oz)

List: \$27.95
FREE delivery Fri, May 9 to Norway on \$49 of eligible items

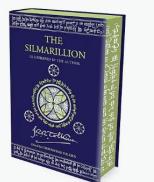
Examines the dark side of internet freedom, written by Evgeny Morozov.



40% off Limited time deal
Open When: A Companion for Life's Twists & Turns



45% off Limited time deal
The Good and Beautiful Bible Study: Experiencing



57% off Limited time deal
The Silmarillion: Illustrated by J.R.R. Tolkien

Problem definition

How to map user queries and LLM output onto the schema to generate useful responses?

At minimum, we need structured LLM output

```
schema product {  
    document product {  
        field make type string {  
            indexing: summary | attribute  
        }  
        field model type string {  
            indexing: summary | attribute  
        }  
        field year type int {  
            indexing: summary | attribute  
        }  
        field price type int {  
            indexing: summary | attribute  
        }  
        field transmission type string {  
            indexing: summary | attribute  
        }  
        field mileage type int {  
            indexing: summary | attribute  
        }  
        field fuelType type string {  
            indexing: summary | attribute  
        }  
        field engineSize type float {  
            indexing: summary | attribute  
        }  
    }  
}
```

Prompt

Prompt - the one used in this demo is 2k

No hallucinations, as our code can trim the JSON output if needed

I'm a user of a car search site. I have some preferences. For example, I might prefer the car to be cheap and I care about the age more about than the mileage.

I need you to transpose these preferences into a JSON that looks like this, so I can use it for queries:

```
{  
  "price": 2.5,  
  "age": 1.7,  
  "mileage": 0.5  
}
```

Possible keys of the JSON are:

- manual/automatic for transmission
- petrol/diesel for fuel fuel type
- age
- price
- mileage
- tax
- mpg
- "make" and "model" of the car

And the number represents the weight: how much I care about the car to be of a specific make, for example. ...

Example LLM output

“I prefer safe european cars”

- Note that the LLM outputs not only brands but other criteria, too, like a moderately new car with low mileage
- This is great - using the collective wisdom encoded into the model

“If I say I want a safe car, you should infer that I care about the age, the mileage and car brands that are known to be safe.”

```
{  
  "age": 1.5,  
  "mileage": 1.0,  
  "volvo": 1.2,  
  "audi": 1.0,  
  "mercedes-benz": 1.0  
}
```

Extract, transform, and load - ETL



Car data

model, year, price, transmission, mileage,

fuelType, tax, mpg, engineSize

Focus, 2016, 8450, Manual, 56022,

Diesel, 0, 74.3, 1.5

Kuga, 2013, 9650, Manual, 44808,

Diesel, 200, 47.9, 2.0

```
{  
  "id": "id:product:product::6ce19adf",  
  "fields": {  
    "mpg": 74.3,  
    "mileage": 56022,  
    "price": 8450,  
    "model": "focus",  
    "year": 2016,  
    "make": "ford",  
    "transmission": "manual",  
    "engineSize": 1.5,  
    "fuelType": "diesel"  
  }  
}
```

Car data, for ranking

Add a “mapped” tensor field for ranking

Normalize feature values, so the user can easily balance feature importance

Some features are binary, like *make* and *transmission* - others continuous, like *price* and *age*.

```
{  
  "id": "id:product:product::6ce19adf",  
  "fields": {  
    "mpg": 74.3,  
    "price": 8450,  
    "model": "focus",  
    "year": 2016,  
    "car_features": {  
      "type": "tensor<float>(features{})",  
      "cells": {  
        "manual": 1.0,  
        "focus": 1.0,  
        "price": 0.9154999852180481,  
        "mpg": 0.7429999709129333,  
        "age": 0.6399999856948853,  
        "ford": 1.0  
      }  
    },  
    "make": "ford",  
    "transmission": "manual",  
    "engineSize": 1.5,  
  }  
}
```

Logstash config - data types and transforms

```
csv {
    columns => [ "model", "year", "price", "transmission", "mileage", "fuelType",
"tax", "mpg", "engineSize" ]
    convert => {
        "year" => "integer"
        "price" => "integer"
        "mileage" => "integer"
        "tax" => "integer"
        "mpg" => "float"
        "engineSize" => "float"
    }
}

# price score: 0.0 to 1.0 based on price, with 100000 as ceiling, 0.0 if above
ceiling
price = event.get('price')
price_score = price <= 100000 ?
(100000 - price) / 100000.0 : 0.0
```

Logstash config, continued

```
# Add numeric preferences only if non-zero
car_features['age'] = age_score if age_score > 0
car_features['price'] = price_score if price_score > 0
car_features['mileage'] = mileage_score if mileage_score > 0
car_features['tax'] = tax_score if tax_score > 0
car_features['mpg'] = mpg_score if mpg_score > 0

# Extract make from filename and add it to preferences
path = event.get('[log][file][path]')
make = File.basename(path, '.csv').split(/\d+/).first.downcase

# add the make as both a proper field and as a preference
event.set('make', make)
car_features[make] = 1.0

# model as a preference, too
model = event.get('model')
car_features[model] = 1.0 if model

event.set('car_features', car_features)
```

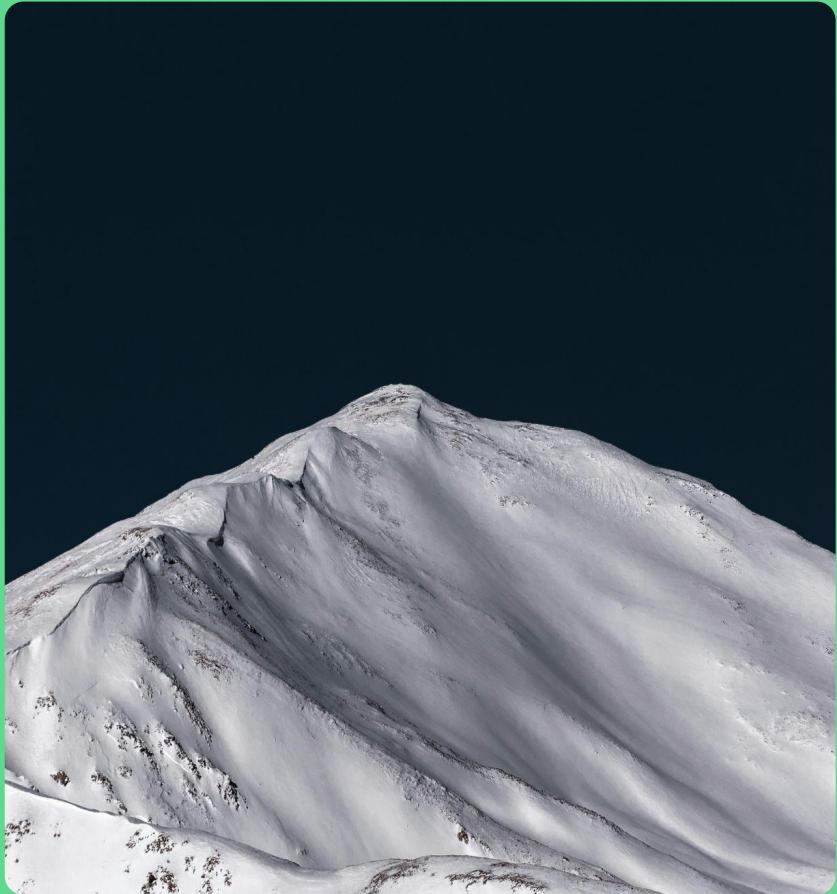
Vespa document

Logstash ETL makes it trivial to generate the ranking features

Note: Keep this simple. You can deploy more math on the values in the ranking profile(s)

```
{  
  "id": "id:product:product::6ce19adf",  
  "fields": {  
    "mpg": 74.3,  
    "price": 8450,  
    "model": "focus",  
    "year": 2016,  
    "car_features": {  
      "type": "tensor<float>(features{})",  
      "cells": {  
        "manual": 1.0,  
        "focus": 1.0,  
        "price": 0.9154999852180481,  
        "mpg": 0.7429999709129333,  
        "age": 0.6399999856948853,  
        "ford": 1.0  
      }  
    },  
    "make": "ford",  
    "transmission": "manual",  
    "engineSize": 1.5,  
  }  
}
```

Queries



Vespa query

```
{  
  "select": {  
    "where": "true"  
  },  
  "ranking": "rank_cars",  
  "ranking.features.query(user_preferences)": "{  
    {features:age}: 2.0,  
    {features:mileage}: 2.0,  
    {features:audi}: 1.5,  
    {features:bmw}: 1.5,  
    {features:mercedes}: 1.5,  
    {features:volkswagen}: 1.5  
  }",  
}
```

Ranking with dot-product

```
rank-profile rank_cars {  
    inputs {  
        query(user_preferences) tensor<float>(features{})  
    }  
  
    first-phase {  
        expression: sum(query(user_preferences) * attribute(car_features))  
    }  
}
```

Example

query(user_preferences)

```
audi,   bmw,   byd, price,   age  
[     2,      3,     -4,      3,      1]
```

attribute(car_features)

```
audi, price, mpg, age, a1  
[1.00, 0.701, -0.547, 0.225, 0.000]
```

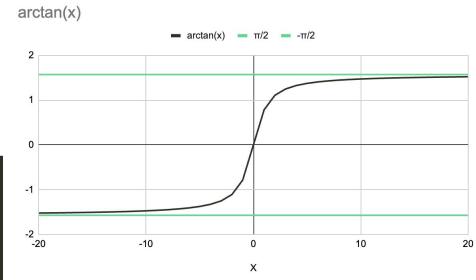
```
audi + price + age  
[2.00 + 2.103 + 0.225 = 4.328]
```

Deploy to prod



Hybrid ranking with normalizations

```
rank-profile rank_cars_hybrid {  
    inputs {  
        query(user_preferences) tensor<float>(features{})  
        query(alpha)  
    }  
    function normalized(val) {  
        expression: atan(val/8) * 2/(3.14159)  
    }  
  
    first-phase {  
        expression {  
            sum(query(user_preferences) * attribute(car_features)) +  
            alpha * normalized(bm25(description))  
        }  
    }  
}
```



Ranking ++

```
rank-profile rank_cars_hybrid {  
  
    inputs {  
        query(user_preferences) tensor<float>(features{})  
    }  
  
    first-phase {  
        expression {  
            sum(query(user_preferences) * attribute(car_features)) +  
            closeness(location).logscale +  
            0.2 * bm25(description) +  
            xgboost("my_model.json")  
        }  
    }  
}
```

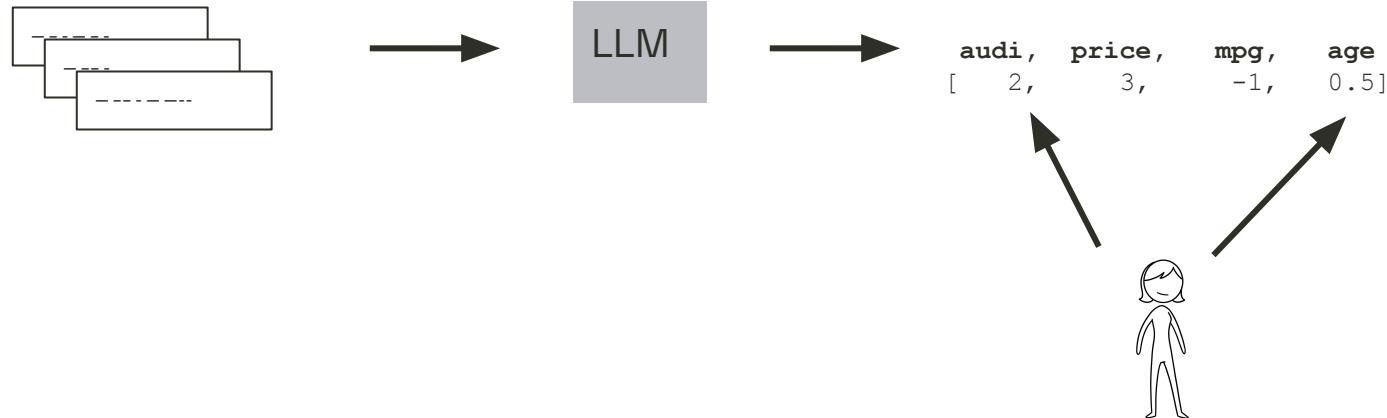
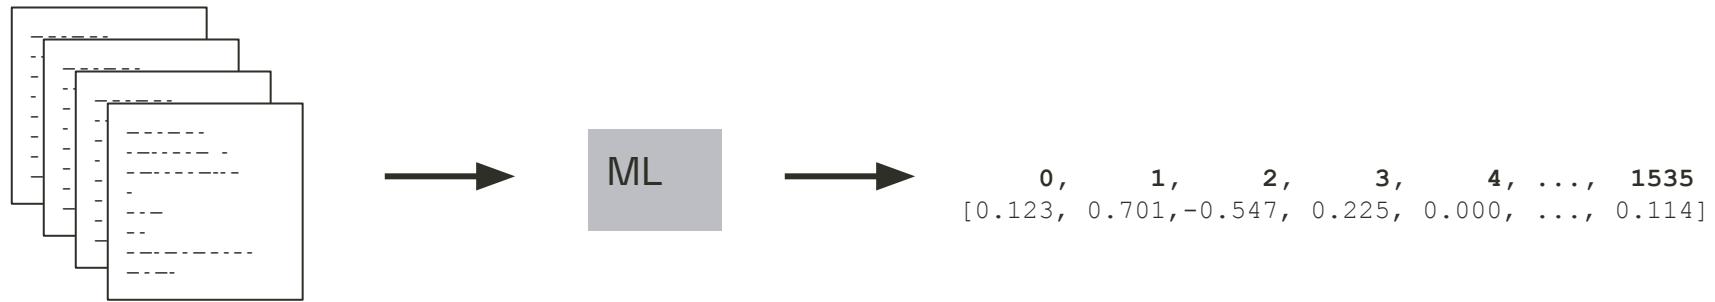
UX

- We are adding a feature where you express preferences in natural language
- Display the preferences, so users can compensate (“I did not like the Audi’s *_that_* much better than BMW ...”)
- We did not explore/build a full app, where, say, one could tune the “preferences” contribution as a whole in the ranking, relative to other signals
- Consider whether to display LLM textual output - not doing so might be a softer start - users type in all sorts of stuff
- How to educate users of the new capability - show some canned questions ?

Tensor ranking



Dense and mapped tensors



Comment



Sparse tensor dot product. First, we define two sparse tensors UserPrefs and Car. Notice the cells that haven't been defined does not have a value:

UserPrefs



Expression `tensor(features{}):{price:0.7, bmw:1.0, audi:1.5}`

Type `tensor(features{})`

Value

features value

audi	1.5
bmw	1
price	0.7

Car



Expression `tensor(features{}):{price:0.66, audi:1.0, mpg:0.7}`

Type `tensor(features{})`

Value

features value

audi	1
mpg	0.7
price	0.66

Comment



In the multiplication, only the matching cells are calculated:

Expression



Expression `UserPrefs * Car`

Type `tensor(features{})`

Value

features value

audi	1.5
price	0.462



Comment

The dot product is the multiplication followed by a sum over the common dimension:

Expression

Expression `sum(UserPrefs*Car, features)`

Type double

Value



value

Comment

This is equivalent to:

Expression

Expression `reduce(join(UserPrefs, Car, f(x,y)(x*y)), sum)`

Type double

Value



value



- Select example ...
- Dense tensor dot product
- Sparse tensor dot product
- Vector-matrix product
- Matrix multiplication
- Tensor generation, dimension renaming and concatenation
- Jaccard similarity between mapped (sparse) tensors
- Neural network


[View setup](#) [Clear](#) [Copy url](#)

Comment

Neural network example (3x5x1). First we define the network input:

input

Expression `tensor(input[3]):[0.1, 0.1, 0.1]`

Type `tensor(input[3])`

Value

input value

0	0.1
1	0.1
2	0.1

Comment

Then we define the hidden layer weights and bias:

whidden

Expression `tensor(input[3],hidden[5]):[[-0.71892,1.99128,0.04418,-0.18243,0.02499], [0.68785,-1.62458,0.23699,0.65774,0.59658], [-1.54464,0.29004,0.16615,1.74138,0.31435]]`

Type `tensor(hidden[5],input[3])`

Value

Try the Tensor playground!

<https://docs.vespa.ai/playground>

<https://docs.vespa.ai/playground/#N4KABGBEBmkFxqNrqmUrWOPYAd50GNIAaFDSPBdDTAF30gGUcBDAJwGcBTMWrqOw5Y2YACZZaYHGyuiArqVoA6MADEAlp1rEwAdx6iu0dfx61dWMB1aczAoZzAADAKrc2ABTZGOTsC35RZwBhdicVAB1+ADkJdQlzAAseBIAbVI5eRJZJbIA3AQBySOAjLgExIxMuIPEuTP4JMHyeFjA8IIS5LjhIMIOAX36B0gxqclwGLhJ+iAp8cZplfgY3Lk9vaA4ZmkxphEg+QWEACmguHLLvDmABqEo4YGI4noAGJQBGHRKAW104D5KV46FhyUTqAFAgZ9XZDDAjqfql0aTA7TUZLSHoWaYFYHUJsHa7KD7KBHBxnC60K71W4PJ5sF5wd4ANmZILBEMBr2h2Lhg3RECRUBRhEJ5ExQrm9AOAEI+FkeD85KlaOocK14jl1Fh+DodakAJ4KsA-HIERImAdmYDSGX83htn0lupyNV6v0GAqxu2wmMqa0x8yo2LmuKgYqWJMqaw2PjAACowPiYTQ+RAEWNAuLIEQvXMJSHSQwACrJMRNaSuBSSdsZWhlpUqtUaghanVgaBYdJYfRBEpGtocOO-MBYAoiespLA-H7t8E-eza-ju2GexFZvOBn0FoIQMOQCPkKNDn4nGNeHzx-E6c6Xa53FPwteZn3Z30BndEw4li2ZWtgLgAEc5HUDpUqESRaCwFdU2fVAhV9KZDugT8fX3ZC5ijbx5ASE4ACssBMM93AvLYdGvDsTgAD2IA07mo+M6LuHQTwfDOnwqABdEABiAA>

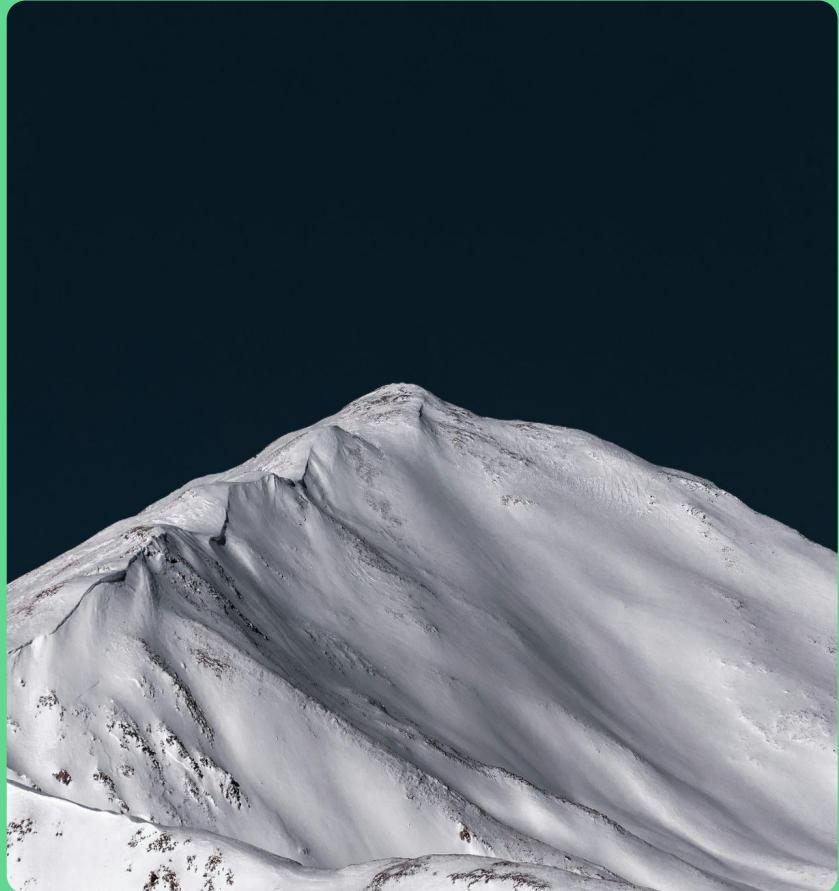
Summary

From natural language to a map(ped tensor) of preferences (features): LLM (+ user input)

Car data: Transform to normalized features using Logstash

Score cars with a dot product over the user prefs and car features

Add this score in a hybrid ranking profile







Follow us!

Vespa.ai

@vespaengine

linkedin.com/company/vespa-ai/

linkedin.com/in/kraune/

