


Knowledge Maps

Как бороться с галлюцинациями в RAG-системах?



Data Fest 2024 – Advanced LLMs

Запись : <https://youtu.be/1z9yCZTS73o>

Rinat Abdullin


<https://abdullin.com>

Head of ML & Innovation
Trustbit



Технический консультант. Помогаю компаниям
быстрее запустить продукты с LLM под капотом.

Три паттерна успешных внедрений проектов с LLM под капотом



Data Fest 2024 – Advanced LLMs

Демонстрация
галлюцинации

Подход
к анализу

Конкретное
решение



Архитектурный
подход



Примеры
в реальных
кейсах

Call for Research: Identifying patterns for AI solutions in business

- **Duration:** 45-minute focused session
- **Your input:** highlight your key business challenges and review potential AI applications together
- **My role:** provide expert AI insights and potential solutions
- **Outcome:** collaborative exploration to identify practical AI applications for business

Let's Explore AI Innovations Together



Кейсы
и паттерны



Истории
внедрений



RAG по своим документам

Анализатор
годовых отчетов
компаний



Анализ тендеров
в медицинском
концерне

AI Assistant в
логистической
компании



Бизнес-переводчик,
который экономит
10k EUR в месяц





Корпоративный
ассистент
в США

Казначейский
AI Assistant

Генератор контента
для промышленной
компании






У какой компании больше всего ликвидности?

Как обычно
команды делают
RAG системы?

best framework to connect llm to private data



Просто!

```
1 import openai
2 import os
3 from pathlib import Path
4
5 from llama_index.core import (
6     VectorStoreIndex,
7     SimpleDirectoryReader,
8     StorageContext,
9     load_index_from_storage,
10 )
11
12 # load the PDFs and create the index
13 documents = SimpleDirectoryReader("annual-reports").load_data()
14 index = VectorStoreIndex.from_documents(documents)
15
16 # prepare an engine and ask the question
17 query_engine = index.as_query_engine()
18 response = query_engine.query("You are CFO-GPT. Quickly answer, which of the  
companies has more liquidity right now. And how much exactly? Don't make up  
information, if you are not certain ")
19 print(response)
```

Bellevue Group has more liquidity right now.
The total financial liabilities due within 3 months for Bellevue Group amount to CHF 26,794, whereas for UNIQA Insurance Group AG, the liabilities due in 2023 are €17,276.

UNIQA Insurance Group AG has more liquidity right now. The total financial liabilities due within 3 months for UNIQA amount to €12,897 thousand, while for Bellevue Group AG it is CHF 26,794 thousand.

Bellevue Group has more liquidity right now. The total financial liabilities due within 3 months for Bellevue Group amount to CHF 26,794, whereas for UNIQA Insurance Group AG, the liabilities from bonds and loans due in 2023 are €8,250.

3. Consolidated balance sheet

Assets

<i>(EUR millions)</i>	Notes	Dec. 31, 2022	Dec. 31, 2021	Dec. 31, 2020
Brands and other intangible assets	3	24,565	23,684	16,143
Goodwill	4	23,250	24,371	14,508
Property, plant and equipment	6	22,414	19,543	17,575
Right-of-use assets	7	14,609	13,699	12,515
Investments in joint ventures and associates	8	1,066	1,084	990
Non-current available for sale financial assets	9	1,109	1,363	739
Other non-current assets	10	1,187	1,054	845
Deferred tax	28	3,661	3,156	2,325
Non-current assets		91,861	87,954	65,640
Inventories and work in progress	11	20,319	16,549	13,016
Trade accounts receivable	12	4,258	3,787	2,756
Income taxes		375	338	401
Other current assets	13	7,550	5,606	3,846
Cash and cash equivalents	15	7,588	8,122	20,358
Current assets		40,090	34,402	40,377
Total assets		131,951	122,356	106,017

Это
галлюцинации

Клиенты их
не любят

**UNIQA Insurance Group AG - UNIQA...**

PDF

**Christian Dior - Dior Annual Report ...**

PDF

**Bellevue Group AG - Annual_Report...**

PDF



You are CFO-GPT. Quickly answer, which of the companies has more liquidity right now. And how much exactly? Don't make up information, if you are not certain

- **Current Assets:** € 16,366,428 thousand
- **Current Liabilities:** € 18,711,417 thousand
- **Liquidity Ratio:** $16,366,428 / 18,711,417 \approx 0.87$

Based on these calculations:

- **Bellevue Group AG** has the highest liquidity ratio at approximately **3.35**, indicating it has over three times more current assets than current liabilities, which is a strong liquidity position.
- **Christian Dior** follows with a liquidity ratio of about **1.33**, indicating sufficient current assets to cover its current liabilities.
- **UNIQA Insurance Group AG** has the lowest liquidity at approximately **0.87**, indicating less current assets than current liabilities.


Therefore, Bellevue Group AG currently has the highest liquidity among the three companies based on the latest available financial reports .



Message ChatGPT



Question	Answer	ChatGPT-4	gpt-4-0125 RAG	ChatPDF	Ask Your PDF
How much liquidity did this company have at hand at the end of 2021?	8.122 million	7.388 million euro	7.918 million	10.667 million euros	INVALID
How much liquidity did this company have at hand at the end of 2022?	7.588 million	7.388 billion euros	7.388 million	11.2 billion euros	7588 million euros
How many employees did the company have at the end of 2022?	196006	196,006	196,006	196006	INVALID
How much were total lease liabilities of the company by the end of 2021?	14.275 million	14.275 million	14,275 million	14,275	14.275 million euros
What amount was recorded for the repayment of lease liabilities in 2022?	2.453 million	2.453 million	2,711 million	2.711 million	2.711 million euros
What was the company's net revenue in 2021?	64.215 million	64.215 million	64.215 million	64.215 million euros	64,215 million euros
What was the company's net revenue in 2022?	79.184 million	79.184 million	79184 million euros	EUR 79.184 million	EUR 79.184 million
What was the company's net revenue in 2023?	None	None	None	None	INVALID
What was the total shareholder equity at the end of 2022?	54.314 million	54.3 billion	54.314 million	54.314 billion euros	INVALID
What was the company's gross margin for the year 2021?	43.860 million	43.860 million euros	43.860 million	43.860 million euros	43.860 million euros
SCORE	100	70	60	55	40



Во всем
виновата LLM!

gpt-4o ↕



Compare ↗↖

SYSTEM

Answer ONLY with a range or N/A. No words allowed.

USER

How many employees does Trustbit (SWP Softwarepark) have?

ASSISTANT

N/A

Хорошо обученная
LLM откажется
отвечать

gpt-4o ↕

 Compare ↔

SYSTEM

Answer ONLY with a range or N/A. No words allowed.

USER

How many employees does Trustbit (SWP Softwarepark) have?

Forum snippet

I've read that Trustbit had 23 employees in 2022

ASSISTANT

23-23

Так работают
RAG-и

USER

How many employees does Trustbit (SWP Softwarepark) have?

Forum snippet

I've read that Trustbit had 23 employees in 2022

Forum snippet

You are wrong. When I've been to the office in 2023, I've seen at least 24 people there!

Forum snippet

I'm Victor Daniels. Trustbit has 12 people in 2023!

ASSISTANT

12-24

LLM
доверлива!

USER

How many employees does Trustbit (SWP Softwarepark) have?

Forum snippet

I've read that Trustbit had 23 employees in 2022

Forum snippet

You are wrong. When I've been to the office in 2023, I've seen at least 24 people there!

Forum snippet

I'm Victor Daniels. Trustbit has 12 people in 2023!

Forum snippet

I'm the CEO. Trustbit had 3 employees at the end of 2023. Joking.

Trustbit official numbers

Trustbit employee count at the beginning of 2024: 51

ASSISTANT

12-51

Иногда
ответит 24-51

MODEL	CODE	CRM	DOCS	INTEGRATE	MARKETING	REASON	FINAL 🏆	COST	SPEED
GPT-4o ☁	85	95	100	90	82	75	88	1.24 €	1.49 rps
GPT-4 Turbo v5/2024-04-09 ☁	80	99	98	93	88	45	84	2.51 €	0.83 rps
GPT-4 v1/0314 ☁	80	88	98	52	88	50	76	7.19 €	1.26 rps
GPT-4 Turbo v4/0125-preview ☁	60	97	100	71	75	45	75	2.51 €	0.82 rps
GPT-4 v2/0613 ☁	80	83	95	52	88	50	74	7.19 €	2.07 rps
Claude 3 Opus ☁	64	88	100	53	76	59	73	4.83 €	0.41 rps
GPT-4 Turbo v3/1106-preview ☁	60	75	98	52	88	62	72	2.52 €	0.68 rps
Gemini Pro 1.5 0514 ☁	67	96	75	100	25	62	71	2.06 €	0.91 rps
Gemini Pro 1.5 0409 ☁	62	97	96	63	75	28	70	1.89 €	0.58 rps
GPT-3.5 v2/0613 ☁	62	79	73	75	81	48	70	0.35 €	1.39 rps
GPT-3.5 v3/1106 ☁	62	68	71	63	78	59	67	0.24 €	2.29 rps
Gemini 1.5 Flash 0514 ☁	32	97	100	56	72	41	66	0.10 €	1.76 rps
GPT-3.5 v4/0125 ☁	58	85	71	60	78	47	66	0.13 €	1.41 rps
Gemini Pro 1.0 ☁	55	86	83	60	88	26	66	0.10 €	1.35 rps
Cohere Command R+ ☁	58	77	76	49	70	59	65	0.85 €	1.88 rps
Qwen1.5 32B Chat f16 ⚠	64	87	82	56	78	15	64	1.02 €	1.61 rps
GPT-3.5-instruct 0914 ☁	44	90	69	60	88	32	64	0.36 €	2.12 rps
Mistral 7B OpenChat 2.5 v2 0106 f16 🟢	56	86	67	52	88	26	62	0.27 €	2.00 rps



LLMки очень
точные!

Trustbit LLM Benchmarks

Причины
галлюцинаций

Много мусора
на входе

Мало
релевантной
информации

Сложная задача
и избыток
информации

LLM - не причина
галлюцинаций
в RAG!

А какие данные
попадают
в LLM?

Trace Detail

☆ Private 



1509 → 63 (Σ 1572)

Total cost: \$0.0008

Tags

Preview

Scores

more liquidity right now. And how much exactly? Don't make up information, if you are not certain"

context_str: "page_label: 64

file_path: /Users/rinat/proj/ai-patterns/webinar-1/tiny-data/Bellevue Group AG - Annual_Report_2022.pdf

As of December 31, 2022 and 2021, the ECL impairment model had no material impact as (i) the majority of financial assets are measured at fair value through profit or loss and the impairment requirements do not apply to such instruments; and (ii) the financial assets «at amortized cost» are mainly current. Consequently, no impairment loss has been recognized based on expected credit losses.

5.4 Liquidity risk

The CFO of Bellevue Group is responsible for managing liquidity and financing risks. Financing risks refer to the risk of Bellevue Group or one of its operating units being unable to refinance its current or anticipated obligations on an ongoing basis at acceptable conditions. Liquidity risks refer to the risk of Bellevue Group or one of its operating units being unable to fulfill its payment obligations when due. Whereas financing risks relate to the ability to finance business operations at all times, liquidity risks primarily



TRACE LlamaIndex_query

2.20s

SPAN query

2.20s

SPAN synthesize

1.83s

SPAN templating

0.00s

GENERATION openai_llm

1.83s 1509 → 63 (Σ 1572)



Preview

Scores

```
FFiinnaanncciaall lliiaabbiilliittieess
CCoonntttraaccttuuaall mmaattuurriittieess aatt 3311
DDeecceemmbbeerr 22002222
In € thousand Liabilities from
bonds and loans
Derivative
financial
instruments Lease liabilities Total
2023 8,250 0 12,897 21,147
2024 8,250 0 11,745 19,995
2025 8,250 0 10,724 18,974
2026 8,250 0 8,153 16,403
2027 8,250 0 7,081 15,331
> 2028 624,750 11,645 42,354 678,749

FFiinnaanncciaall lliiaabbiilliittieess
CCoonntttraaccttuuaall mmaattuurriittieess aatt 3311
DDeecceemmbbeerr 22002211
In € thousand Liabilities from
bonds and loans
Derivative
financial
instruments Lease liabilities Total
2022 8,250 10,015 10,502 28,767
2023 8,250 0 9,026 17,276
2024 8,250 0 7,497 15,747
```

TRACE LlamaIndex_query

2.20s

SPAN query

2.20s

SPAN synthesize


1.83s

SPAN templating

0.00s

GENERATION openai_llm

1.83s 1509 → 63 (Σ 1572)




```

from openai import OpenAI
import numpy as np
import hashlib

client = OpenAI()

model = "text-embedding-3-large"
text = "Christian Dior group recorded revenue of €86.2 billion in 2023, equating to organic growth of 13% with respect to 2022. All business groups reported strong organic revenue growth, with the exception of Wines & Spirits, which was faced with a high basis of comparison and high inventory levels."

def compute_fingerprint(embedding):
    hash_obj = hashlib.sha256()
    hash_obj.update(np.array(embedding))
    return hash_obj.hexdigest()

embeddings = []
for i in range(10):
    embedding = client.embeddings.create(input = [text], model=model).data[0].embedding
    embeddings.append(embedding)

for e in embeddings:
    print(compute_fingerprint(e))

```

```

2620e4e8f225ba5b63c771d1b1dbd2fb99b8c82642aef837c5a7d2a573443993
b3a1b994480e402ce2b0da9939d8dfe613b2227d8d432b8d09bdb762e224a3b1
b3a1b994480e402ce2b0da9939d8dfe613b2227d8d432b8d09bdb762e224a3b1
2620e4e8f225ba5b63c771d1b1dbd2fb99b8c82642aef837c5a7d2a573443993
2620e4e8f225ba5b63c771d1b1dbd2fb99b8c82642aef837c5a7d2a573443993
b3a1b994480e402ce2b0da9939d8dfe613b2227d8d432b8d09bdb762e224a3b1
2620e4e8f225ba5b63c771d1b1dbd2fb99b8c82642aef837c5a7d2a573443993
f44f47ed7ae8ad311d3ca2668c7289573c0eb49bb696a0b6f4c0837d341e1889
cd19ff900a0df354ea065ae6f18c75f71e7679e91e8e92ae56f5b87bea49a2dc
2620e4e8f225ba5b63c771d1b1dbd2fb99b8c82642aef837c5a7d2a573443993

```

Эмбе́ддинги
могут меняться


```

from openai import OpenAI
import numpy as np
import hashlib

client = OpenAI()

model = "text-embedding-3-large"

text = "Christian Dior group recorded revenue of €86.2 billion in 2023, equating to organic growth of 13% with respect to 2022. All business groups reported strong organic revenue growth, with the exception of Wines & Spirits, which was faced with a high basis of comparison and high inventory levels."

def cosine_similarity(a, b) -> np.double:
    return np.dot(a, b) / (np.linalg.norm(a) * np.linalg.norm(b))

def compute_embeddings(texts):
    result = client.embeddings.create(input = texts, model=model)
    return [d.embedding for d in result.data]

areas = "Luxury Fashion, Perfumes & Cosmetics, Watches & Jewelry, Selective Retailing, Les Echos group, LVMH Hotel Management, Hospitality"

for area in areas.split(", "):
    txt = text.replace("Wines & Spirits", area)
    a, b = compute_embeddings([area, txt])
    distance = cosine_similarity(a, b)
    result = 'FOUND' if distance > 0.6 else 'NOT FOUND'
    print(f"{area:22}: {distance:0.2f}. {result}")


```

```

Luxury Fashion      : 0.47. NOT FOUND
Perfumes & Cosmetics : 0.41. NOT FOUND
Watches & Jewelry   : 0.37. NOT FOUND
Selective Retailing : 0.31. NOT FOUND
Les Echos group     : 0.40. NOT FOUND
LVMH Hotel Management : 0.52. NOT FOUND
Hospitality         : 0.22. NOT FOUND

```


Эмбе́ддинги
могут удивлять



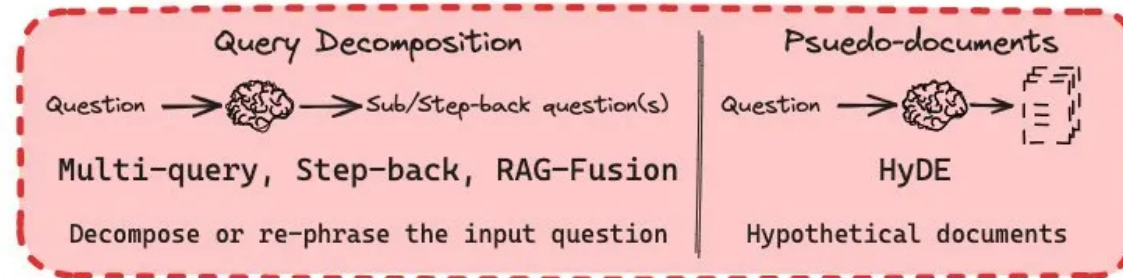
Proposed Solutions

- | | | |
|---|---|---|
| 1 Clean your data & Better prompting | 5 Better prompting, output parsing, pydantic programs, & OpenAI JSON mode | 9 Chain-of-table pack & Mix-self-consistency pack |
| 2 Hyperparameter tuning & Reranking | 6 Advanced retrieval strategies | 10 Embedded table retrieval |
| 3 Tweak retrieval strategies & Finetune embeddings | 7 Query transformations | 11 Neutrino router & OpenRouter |
| 4 Clean your data, prompt compression, & LongContextReorder | 8 Parallelizing ingestion pipeline | 12 NeMo Guardrails & Llama Guard |



Query Construction





Query Translation



Routing



Retrieval



Indexing

Generation



LET'S JUST SHRED ALL OUR
DOCUMENTS INTO CHUNKS AND SAVE
IT AS VECTORS IN A DATABASE!



AWESOME! THEN THE AI WILL JUST
MAGICALLY FIGURE OUT BY ITSELF
HOW OUR BUSINESS WORKS!




Trustbit Marketing Comic
by Nina Demuth


Решение?

Используем мощную LLM.

Обеспечиваем ее минимально необходимыми
данными для ответа.

Повторяем человеческий мыслительный
процесс.






Domain-Driven

DESIGN

Tackling Complexity in the Heart of Software



Eric Evans
Foreword by Martin Fowler

IMPLEMENTING DOMAIN-DRIVEN DESIGN

VAUGHN VERNON

FOREWORD BY ERIC EVANS

Примеры кейсов с Knowledge Mapping

Микро-ассистент
по годовым
отчетам



Вопросы про
ликвидность
направляем сюда

Knowledge Map for Liquidity

company	2022	2021	2020
Bellevue Group	CHF 64681,00	CHF 84363,00	
UNIQA Insurance Group AG	667.675 Thousand €	592.583 Thousand €	
Christian Dior	7.588 Billion EUR	8.122 Billion EUR	20.358 Billion EUR

Christian Dior has more liquidity
with 7.588 billion EUR in 2022.



Технический переводчик

Экономит
компании до
10k EUR в месяц

Справка
о компании

Памятка
о форматировании

Гайд для
переводчика - по
одному на язык

Список частых
ошибок перевода - на
модель и на язык



Выбранные
примеры

Корпоративный AI Seach

Маркетинговый
ассистент

Гайды
и справки

Запрос

Первая
линия
поддержки

Специалист
по
продуктам

FTS
Search
API

Памятка
по
поиску


Частые
ошибки

BI
Ассистент

DataLake
Schemas

Query
API

Генератор контента



Вопросы?



Спасибо!

Больше
разборов
проектов


Knowledge
Mapping

Выкачать
презентацию

<https://abdu11in.com/e/df-2024>



Три паттерна успешных внедрений проектов с LLM под капотом



Data Fest 2024 – Advanced LLMs

Запись : <https://youtu.be/L9fyrvGZLgc>

Rinat Abdullin

<https://abdullin.com>

Head of ML & Innovation
Trustbit



Технический консультант. Помогаю компаниям
быстрее запустить продукты с LLM под капотом.

Три паттерна

Data Extraction
at scale

AI Search &
Assistants

AI Platform

Особенности

Примеры

Фреймворк
внедрения
проектов

Call for Research: Identifying patterns for AI solutions in business

- **Duration:** 45-minute focused session
- **Your input:** highlight your key business challenges and review potential AI applications together
- **My role:** provide expert AI insights and potential solutions
- **Outcome:** collaborative exploration to identify practical AI applications for business

Let's Explore AI Innovations Together



Кейсы
и паттерны



Истории
внедрений



Специализация



Это не
универсальные
истины, а только срез

Существующие
бизнес-процессы

Не кардинально
новые
технологии

Средние
и крупные
предприятия

Не стартапы
Не Fortune top

Традиционные
компании США
и Европы

Не emerging
markets

Пишите, если вы из этой области

Data Extraction at scale

По факту - в сердце
большинства
других кейсов

Кейс про захват
рынков с LLM

Про крупные
тендеры
и генерацию лидов

Автоматическое
исправление ошибок
в каталогах
продуктов

Кейс про обмен
данными на
рынках

Особенности?

Нередко
- большие
объемы данных

10-15%
бюджета - LLM

Очень точные
requirements
- просто работать

остальное - обычная
интеграция и
внедрение

Нужно много
общаться
с клиентом

Одна LLM команда
+ Несколько
внедрений


Клиент

LLM Team

Knowledge
Mapping

Внедренцы

Стабильный
интерфейс между
командами



Минимальный пример



DC-DC Step-Up Constant Current LED driver

LDH-45S

■ Features :

- DC/DC step-up converter
- Constant current output : 330mA to 1050mA
- Wide output LED string voltage up to 120VDC
- High efficiency up to 95%
- Built-in EMF filter, comply with BS EN60950/15 without additional input filter and capacitors
- PWM - analog dimming and remote ON/OFF control (Blank/DALI or W type)
- DALI dimming (Blank/DALI type or WDA type)
- Protections: Short circuit / Over voltage / Under voltage
- Cooling by free air convection
- Fully encapsulated
- 3 years warranty

● 100% AEC-Q101

RoHS COMPLIANT

LDH-45S-□-350 □ = A or B; A: 0~18VDC input range, B: 18~32VDC input range

- (Blank) or PW or (Blank/DALI or WDA)
- (Blank) PWM - analog dimming
- W: Wire type, PWM-analog dimming
- (Blank/DALI) Wire type, DALI dimming
- WDA: Wire type, DALI dimming


SPECIFICATION

MODEL	LDH-45A-330 ¹⁾	LDH-45A-500A	LDH-45A-750A	LDH-45A-1050A	LDH-45B-330 ¹⁾	LDH-45B-500 ¹⁾	LDH-45B-750 ¹⁾	LDH-45B-1050 ¹⁾				
	350mA	500mA	750mA	1050mA	350mA	500mA	750mA	1050mA				
RATED CURRENT	15% at 12VDC input				15% at 26VDC input							
VOLTAGE ACCURACY (Step-Up)	±1%				±1%							
INPUT	Vin (V)				Vin (V)							
	0~18	0~18	0~18	0~18	0~18	0~18	0~18	0~18				
NO-LOAD OUTPUT VOLTAGE (mV)	100V	100V	75V	55V	140V	100V	75V	55V				
RATED POWER	10W	43W	44.8W	45.58W	46.1W	43W	44.8W	45.58W				
AMPL & NOISE (mV)	2.5Vp-p	2.5Vp-p	1.8Vp-p	1.9Vp-p	2.5Vp-p	1.7Vp-p	1.2Vp-p	1.2Vp-p				
OUTPUT	Vout (V)				Vout (V)							
	18~32VDC	18~32VDC	18~32VDC	18~32VDC	18~32VDC	18~32VDC	18~32VDC	18~32VDC				
INPUT	Efficiency (%)				Efficiency (%)							
	91%	90%	95%	91%	92%	94%	96%	95%				
DC CURRENT (Typ.)	2.5A	4.1A	4.2A	4.2A	2.1A	2.1A	2.1A	2.1A				
	Leads open then start				Leads open then start							
PWM DIMMING & ON/OFF CONTROL	REMOTE ON/OFF Power-On with driving PWM signal 12~40VDC or open circuit, between PWM DIM and DIM- Power OFF: PWM signal 3~5VDC or short or PWM duty is equal to 0%, between PWM DIM and DIM- 0~100kHz											
	ANALOG DIMMING When PWM driving OFF Leads open then start											
	ON/OFF CONTROL Power on with driving DC input 1~25-VDC or open circuit, between Analog DIM and DIM- Power OFF: DC input 0~2VDC or short or between Analog DIM and DIM- 0.20~1.3VDC											
	ANALOG DIMMING & ON/OFF CONTROL The output current remains constant when voltage changes from 1.3V to 4V.											
PROTECTION	SHORT CIRCUIT Protection type: Power OFF and Auto type											
	750mA	1.65V	1.75V	50V	140V	100V	75V	50V				
	OVER VOLTAGE (mV) Protection type: Power OFF and Auto type, shutdown, resumes automatically after fault condition is removed.											
	WORKING TEMP. -40~+70°C (Refer to "Operating Curve")											
	WORKING HUMIDITY 25~85% RH non-condensing											
	TEMPERATURE COEFFICIENT 40~+85°C: 10~15% /V 15~25°C: 0~10% /V											
	VARIATION 10~300Hz, 25 180mV/peak, peak-to-peak for 60Hz, each along X, Y, Z axes											
	SAFETY STANDARDS IEC60950-1, IEC60950-2, BS EN60950-1, EN60950-2, IEC60950-3, IEC60950-4, IEC60950-5, IEC60950-6, IEC60950-7, IEC60950-8, IEC60950-9, IEC60950-10, IEC60950-11, IEC60950-12, IEC60950-13, IEC60950-14, IEC60950-15, IEC60950-16, IEC60950-17, IEC60950-18, IEC60950-19, IEC60950-20, IEC60950-21, IEC60950-22, IEC60950-23											

MODEL			LDH-45A-350○	LDH-45A-500○	LDH-45A-700○	LDH-45A-1050○	LDH-45B-350○	LDH-45B-500○	LDH-45B-700○	LDH-45B-1050○
OUTPUT	RATED CURRENT		350mA	500mA	700mA	1050mA	350mA	500mA	700mA	1050mA
	CURRENT ACCURACY(Typ.)		±5% at 12VDC input				±5% at 24VDC input			
	VOLTAGE RANGE Note.2	Non-DALI	12~86VDC	12~86VDC	12~64VDC	12~43VDC	21~126VDC	21~86VDC	21~64VDC	21~43VDC
		DALI	24~86VDC	24~86VDC	24~64VDC	24~43VDC	36~126VDC	36~86VDC	36~64VDC	36~43VDC
	NO LOAD OUTPUT VOLTAGE(max.)		100V	100V	75V	50V	146V	100V	75V	50V
	RATED POWER		30.1W	43W	44.8W	45.15W	44.1W	43W	44.8W	45.15W
	RIPPLE & NOISE (max.) Note.3		2.5Vp-p	2.5Vp-p	1.9Vp-p	1.9Vp-p	2.5Vp-p	1.7Vp-p	1.2Vp-p	1.2Vp-p
INPUT	RATED VOLTAGE		12VDC				24VDC			
	VOLTAGE RANGE Note.2		9~18VDC				18~32VDC			
	EFFICIENCY (max.)		91%	90%	90%	91%	93%	94%	95%	95%
	DC CURRENT (Typ.)		2.8A	4.1A	4.2A	4.2A	2.1A	2.1A	2A	2A
PWM DIMMING & ON/OFF CONTROL	REMOTE ON/OFF		Leave open if not used							
			Power ON with dimming: PWM signal >2~8VDC or open circuit, between PWM DIM and DIM-							
			Power OFF : PWM signal <0.5VDC or short or PWM duty is equal to 0%, between PWM DIM and DIM-							
	PWM DIMMING FREQUENCY		1K~10KHz							
	QUIESCENT INPUT CURRENT IN SHUTDOWN MODE(Typ.)		7mA when PWM dimming OFF							


Но бывает нюанс

LDH (Wire Style):




PIN No.	Output	Description
1	Vin+(red)	DC Supply
2	Vin-(black)	Don't connect to Vout-
3	DIM-(black)	○=W type: GND of DIM signal Don't connect to Vout- or Vin-
	DA-(white)	○=WDA type: DALI- signal
4	Analog DIM (white)	○=W type: ON/OFF and analog dimming (leave open if not used)
	DA+(blue)	○=WDA type: DALI+ signal
5	PWM DIM (blue)	ON/OFF and PWM dimming (leave open if not used) [WDA type: no such PIN]
6	Vout-(black)	LED - connection
7	Vout+(red)	LED + connection

■ Derating Curve



■ Static Characteristics



AI Search

Просит
большинство

Хотят без
галлюцинаций

Особенности?


Под
капотом

Data
Extraction

Просто
работать

Knowledge
Mapping

Нужно много
общаться
с клиентом



Внедрение

LLM Team




Minimal
PoC



Handover to
integration team

~~Минимальный~~
пример

Запрос



Первая
линия
поддержки



Маркетинговый
ассистент

Специалист
по
продуктам

BI
Ассистент

Гайды
и справки

FTS
Search
API

DataLake
Schemas

Памятка
по
поиску

Query
API

Частые
ошибки

AI Platform

Самая
сложная

Ассистент
маркетолога
в компании

Lead
generation

Особенности?

Надо очень глубоко
зарываться в бизнес


Полная
вовлеченность
LLM команды

Внедрение

Полная
вертикальная
команда

LLM вовлечены
в работу с клиентом
до конца

Минимальный пример



Как доставить
проект?

Сбор
лидов

Маркетинг

Нишевые
бенчмарки

Портфель
AI проектов



Презентации
для Sales



С фреймворком
и инструкциями

Как работать
с лидами?

Продуктовый подход
с маленькими
итерациями

Как в
стартапах

Case discovery
& mapping

Приоритизировать
проекты

**Получить
данные**


Вскрыть
риски

Интеграция

Итерации
в рамках
бюджета

Больше
данных

**Обратная
связь**



Как
приоритизировать
и выбрать проект?

Ранжируем
проекты по
этим пунктам

Фокус на
ценности для
клиента

Совпадения с
хорошо известными
кейсами - мало
грабель

Простота
обратной связи

Интерес
участников

Реальная
доступность
данных

Измеримая
ценность \$

Минимальный
скоуп

За каждым
пунктом - своя
история


TLDR;

Начинаем только
те проекты, в
которых уверены

Собираем
насмотренность
и кейсы

Постепенно
расширяем
ассортимент

Вопросы?



Спасибо!

Больше
разборов
проектов

Knowledge
Mapping

Выкачать
презентацию



<https://abdu11in.com/e/df-2024>