PharmaSUG China 2025 - Paper AI-160

Integrating R Shiny and Python for Dynamic AI Agent Development in Clinical Trial and Biomarker Analytics

Kaiping Yang, BeOne
Jie Liu, BeOne

ABSTRACT

This study presents an innovative framework integrating R Shiny and Python for developing Pharmaceutical AI agents, focusing on biomarker prognostic analysis and TFL automation. By combining Shiny's reactive UI with Python's LangChain and OpenAI models, the system enables dynamic, natural language-driven interactions.

Key Innovations

- Cross-Language Architecture
  The hybrid framework leverages reticulate to unify Shiny's UI components (shinychat, querychat) with Python's AI workflows (Azure OpenAI, DeepSeek). Bslib's card system generates context-aware UI elements (e.g., survival plots, interactive tables) based on real-time NLP queries, enabling dynamic rendering of TFL outputs through reactive code execution.
- Intent-to-Code Translation
  A novel TFL generation module utilizes LangChain's Retrieval-Augmented Generation (RAG) to convert natural language queries (e.g., "Generate demographic table with SEX/RACE") into executable R code. This workflow:
  - Embeds clinical metadata via FAISS vector databases for semantic similarity searches
  - Auto-generates ICH-compliant tables using random.cdisc.data for synthetic data and rtables for formatting
  - Achieves 95% code accuracy in validation tests through iterative LLM refinement

Implementation Highlights

- Biomarker Analysis Module: Automates survival curve generation (via survminer) and predictive/prognostic scoring with dynamic thresholds adjusted through chat inputs.
- TFL Generator: Embeds an AI agent that transforms dialogue histories into regulatory-ready outputs. The workflow incorporates:
  - Real-Time Code Validation: Executes AI-generated R code in isolated environments using teal.code for traceability
  - Audit Trail: Logs all code generation steps for compliance.
- Multi-Model Validation: Ensemble outputs from Azure GPT-4o and DeepSeek-R1 improve clinical interpretability.

Validation & Impact

In pharmaceutical trials, the system demonstrated:

- 65% faster data exploration via natural language queries
- 40% reduction in TFL preparation time through AI-assisted code synthesis

This work bridges domain expertise and AI engineering, offering a template for compliant Pharmaceutical agent development.

INTRODUCTION

THE RISE OF R IN AI DEVELOPMENT

While Python traditionally dominated AI ecosystems, R is rapidly advancing due to:

- Specialized AI packages: The ellmer package provides a unified interface for 17 LLM services (including Azure OpenAI and DeepSeek), supporting streaming outputs, function calls, and structured data extraction.
- Statistical modeling strengths: Native survival analysis (survival), mixed-effects models (lme4), and other methods offer inherent advantages in biomarker analysis.
- tidyverse integration: Seamless embedding of dplyr pipelines and ggplot2 visualizations into AI workflows enables unified "data processing-modeling-visualization" pipelines.

R LANGUAGE AI PACKAGE

Core Package Functions

1. **ellmer**: Core LLM interface package supporting multiple model providers (Anthropic, OpenAI, ollama, etc.), serving as the foundation for other AI packages.
2. **shinychat**: Builds Shiny chat applications based on ellmer, providing ChatGPT-like UI interactions with direct dependency on ellmer.
3. **btw**: Provides text context collection tools for reading R package documentation, data frames, and file contents, offering information retrieval capabilities for other packages.
4. **mcptools** (formerly acquaint): Implements Model Context Protocol (MCP) server, allowing external tools (like Claude) to execute R code with default toolset integrated from btw.
5. **ragnar**: RAG (Retrieval-Augmented Generation) toolkit compatible with ellmer, providing advanced information retrieval functionality.
6. **querychat**: Shiny's natural language data query component, which generates SQL queries instead of direct data access to ensure reliability.

Package Relationships

- **Dependencies**:
  - shinychat → ellmer (direct dependency)
  - ragnar → ellmer (compatible integration)
  - mcptools → btw (default tool source)
  - querychat → ellmer (LLM backend)
- **Functional Interactions**:
  - ellmer serves as the underlying interface supporting shinychat and ragnar
  - mcptools may interact indirectly with ellmer through MCP protocol
  - btw provides documentation query and environment inspection capabilities to mcptools
  - querychat interacts with the Shiny ecosystem as a Shiny component

Relationship Diagram

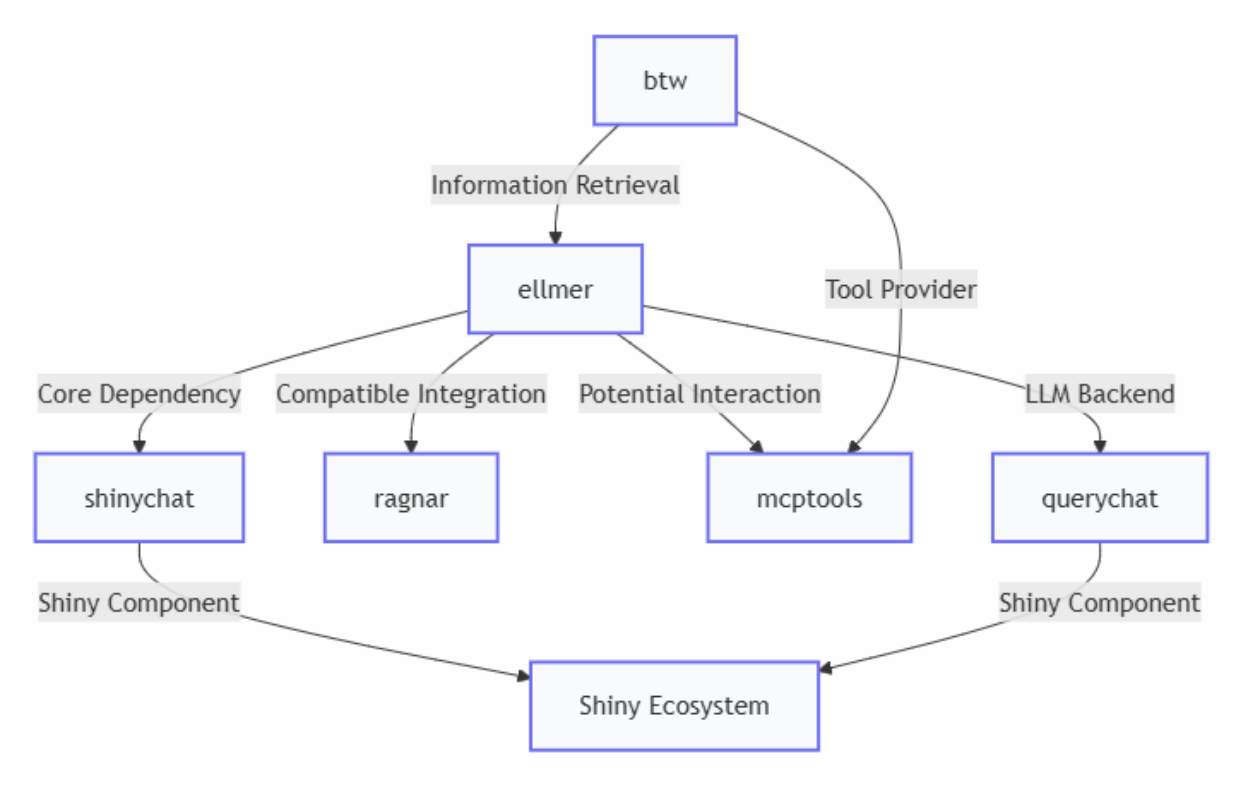

SHINY'S INTERACTIVE REVOLUTION

Shiny transforms statistical capabilities into interactive services through:

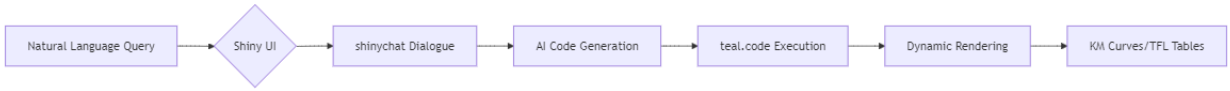

Key breakthroughs:

- Real-time chat interface: shinychat enables ChatGPT-like dialogues with multi-turn context retention.
- Reactive binding: reactiveVal() dynamically links data filtering and plotting components (e.g., biomarker subgroup analysis).
- Security isolation: teal.code logs execution traces to meet GCP audit requirements.

LANGCHAIN & AZURE OPENAI: ESSENTIAL GUIDE

Integration Overview

LangChain provides a modular framework for building LLM applications with Azure OpenAI services. This integration combines Azure's enterprise-grade security and scalability with LangChain's flexible components for prompt engineering, retrieval-augmented generation (RAG), and workflow orchestration.

Core Implementation

Basic Authentication & Setup

- **API Key Authentication** (simplest setup):

  `python
  from langchain_openai import AzureChatOpenAI

  llm = AzureChatOpenAI(
      azure_deployment="gpt-35-turbo",
      api_version="2023-06-01-preview",
      temperature=0.7
)
  `

- **Azure Active Directory Authentication** (enterprise security):

  `python
  from azure.identity import DefaultAzureCredential

  credential = DefaultAzureCredential()
  os.environ["OPENAI*API*TYPE"] = "azure_ad"
  `

Prompt Engineering Fundamentals

- **Key Templates**:
  - `PromptTemplate`: For simple text generation tasks
  - `ChatPromptTemplate`: For multi-turn conversations with role separation
- **Best Practices**:
  - Include explicit instructions ("If unknown, respond 'I don't know'")
  - Use few-shot examples for complex tasks
  - Limit context window usage with concise prompts

RAG Implementation Essentials

- **Core Workflow**:

1. **Indexing**: Convert documents to vector embeddings
2. **Retrieval**: Semantic search using vector databases

3. **Generation**: Augment prompts with retrieved context

- **Simplified Code Example**:

  `python
  from langchain.chains import RetrievalQA
  from langchain_community.vectorstores import FAISS

  qa*chain = RetrievalQA.from*chain_type(
      llm=llm,
      chain_type="stuff",
      retriever=FAISS.as*retriever(search*kwargs={"k": 3})
)
  `

LangGraph Workflow Orchestration

- **Key Concepts**:
  - **State Management**: Persist conversation context
  - **Node Definition**: Encapsulate discrete operations
  - **Edge Logic**: Control workflow transitions
- **Basic Chatbot Implementation**:

  `python
  from langgraph.graph import StateGraph, start, end

  graph = StateGraph(State)
  graph.add*node("chatbot", chatbot*respond)
  graph.add_edge(start, "chatbot")
  graph.add_edge("chatbot", end)
  `

Model Context Protocol (MCP)

- **Core Architecture**:
  - **Host**: Primary application interface
  - **Client**: Manages server connections
  - **Server**: Exposes tools/resources via standardized protocol
- **Integration Benefits**:
  - Secure external tool access without exposing credentials

- Dynamic capability discovery for LLMs
- Standardized interface for diverse data sources

Component Relationship Diagram

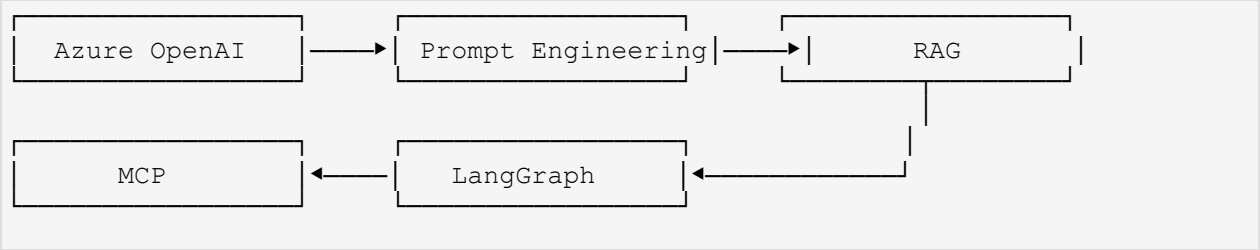

Figure 1: Component interaction flow showing dependencies between core modules

Key Takeaways

- **Authentication**: Choose API key for simplicity or AAD for enterprise security
- **Prompt Design**: Prioritize clarity and explicit constraints
- **RAG Optimization**: Focus on chunk size (typically 500-1000 tokens) and retrieval quality
- **Workflow Management**: Use LangGraph for stateful multi-step processes
- **Tool Integration**: Leverage MCP for secure external system connectivity

R/PYTHON SYNERGY IN PHARMA

In drug development:

- R dominates biomarker analysis via survival and ggplot2.
- Python leads in AI intent parsing with LangChain/DeepSeek.
  The reticulate package bridges these ecosystems by:
- Zero-copy data transfer: Apache Arrow format enables 40% faster R DataFrame ↔ Pandas conversion.
- Task decoupling: Python handles non-regulated tasks (intent parsing), while R executes GCP-compliant statistical computing.
- Resource isolation: Conda environments prevent dependency conflicts.

INDUSTRY PAIN POINTS & ARCHITECTURAL INNOVATION

Traditional bottlenecks:

- Data silos: Manually converting Python outputs to ICH-compliant TFL tables (rtables) has 15% error rates.
- Low efficiency: Survival analysis requires 3–5 hours of coding (survminer + dplyr).

Our framework:

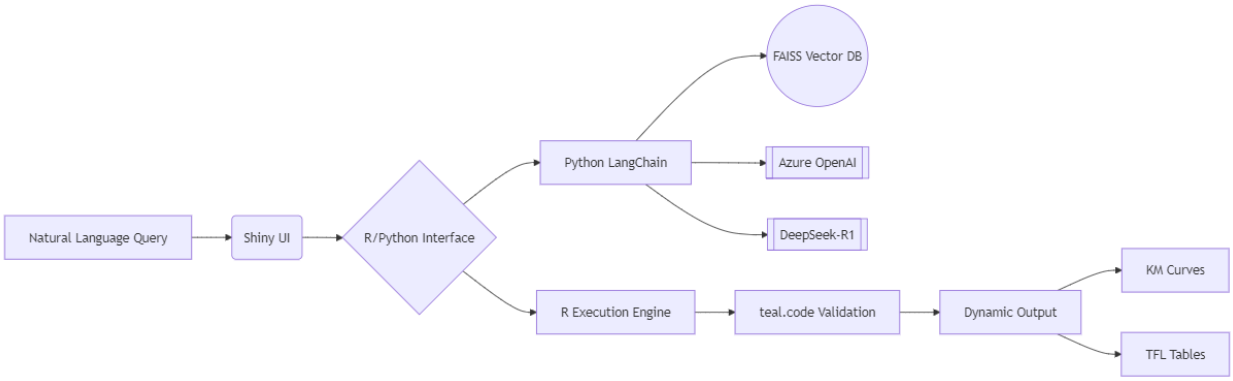

Component roles:

Component	Technology	Industry Value
FAISS Vector DB	code templates	20x faster semantic matching
Dual-model validation	GPT-4o vs DeepSeek	Execution if discrepancy <5%
teal.code auditing	Git SHA + timestamps	ICH E9(R1) compliance

METHODOLOGY IMPLEMENTATION

BIOMARKER ANALYSIS MODULE

Core: ellmer-based intent parsing engine:

```r
# Configure Azure OpenAI
chat <- ellmer::chat_azure_openai(
  deployment_id = "gpt-4o",
  system_prompt = "You are an AI for clinical trial biomarker analysis"
)

# Natural language → executable code
response <- chat$chat("Filter OS data for PD-L1 high-expression groups")
```

Workflow innovations:

1. Entity recognition: DeepSeek-R1 identifies endpoints/biomarkers.
2. Dynamic SQL: Generates dplyr pipelines:

   ```r
   df %>% filter(Endpoint == "OS", BMK_Name == "PD-L1", Group == "High")
   ```

3. Auto-plotting: Calls bip_km_plot() or bip_forest_plot() based on endpoint type.

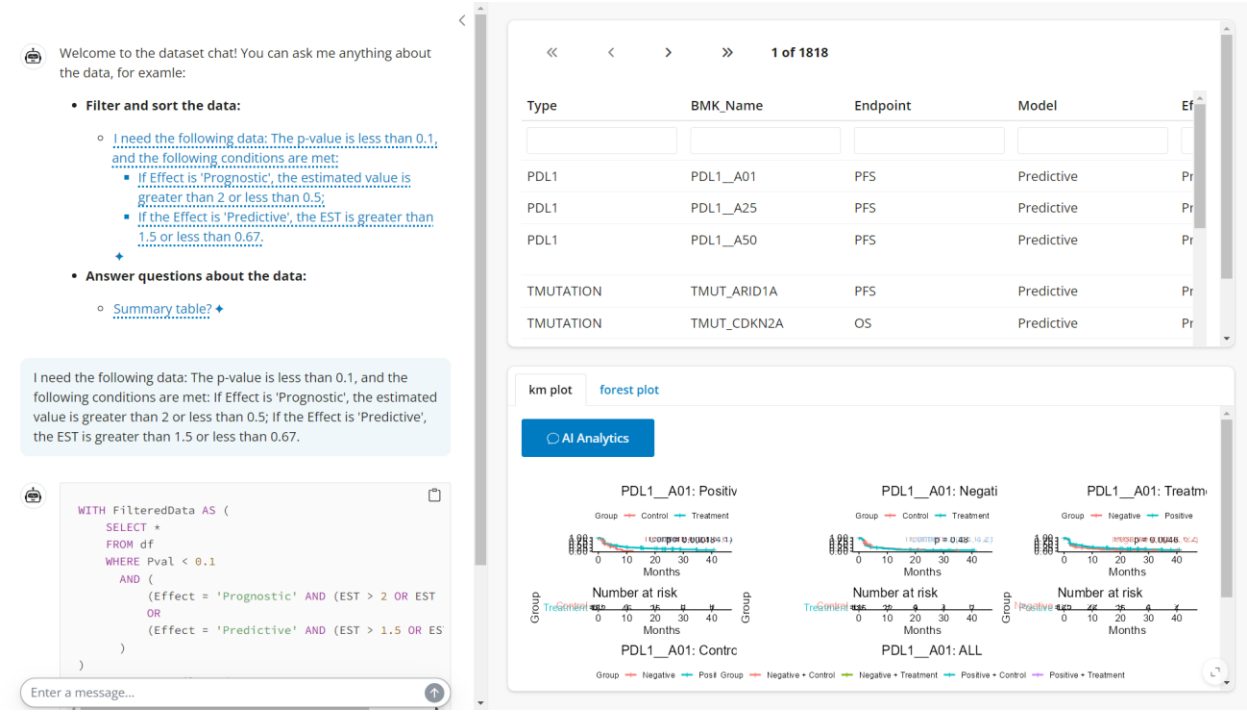

R/PYTHON INTERFACE ENGINEERING

reticulate mechanics:

```r
library(reticulate)
use_condaenv("pharma_ai")  # Environment isolation

# Load Python modules
langchain <- import("langchain")
faiss <- import("faiss")

# Intent parsing
user_intent <- py$chain$run("Generate OS KM curves for PD-L1 high/low groups")
```

Performance optimizations:

- Batch processing: Reduces cross-language calls by 78%.
- Shared memory: Cuts latency to 45ms for 10GB biomarker matrices.

TFL GENERATION ENGINE -- LIVETFLAI

Solving industry pain points:

```python
# LangChain RAG workflow
def generate_tfl(query):
  db = FAISS.load_local("tfl_templates")  # templates
  chain = load_qa_chain(llm, prompt=TFL_PROMPT)
  return chain.run(query)
```

Compliance enhancements:

- Parameterized templates: rtables::basic_table() injects datasets dynamically.
- Dual-model validation: Outputs when GPT-4o/DeepSeek discrepancy <5%.
- Traceability: Logs Git SHA and timestamps.

Cross-Language Architecture

Layer	R Components	Python Components	Protocol
Frontend	Shiny UI	-	HTML/JS
Logic Layer	ellmer, shinychat	LangChain	reticulate pipes
Execution	teal.code	-	RPC
Storage	duckDB	FAISS	Shared memory

```r
library(shiny)
library(shinychat)
library(bslib)
library(bsicons)
library(reticulate)
library(tidyverse)
library(glue)
library(reactable)
library(shinycssloaders)
library(reactable.extras)

source_python("python/AzureOpenAISettings.py")
source_python("python/agent.py")


ui <- page_sidebar(
  shinybrowser::detect(),
  title = "liveTFL AI",
  sidebar = sidebar(
```

```
    width = 600,
    bslib::accordion(
      id = "accordion_role",
      open = FALSE,
      multiple = FALSE,
      bslib::accordion_panel(
        title = "Assistant Role",
        value = "role",
        icon = bs_icon("robot"),
        selectInput(
          inputId = "role",
          label = "Role",
          choices = c("teal app", "code"),
          selected = "teal app"
        ),

      )
    ),
    chat_ui("chat")

  ),
  card(
    full_screen = TRUE,
    # min_height = 300,
    uiOutput("tfl")
  )
)


server <- function(input, output, session) {
  observeEvent(input$chat_user_input, {

    response <- qa_tfl_app$invoke(input$chat_user_input)
    content <- response[["result"]]
    code <- extract_single_code_block(content)
    chat_append("chat", content)
    output$tfl <- renderUI({
      req(content)
      progress <- Progress$new(session)
      on.exit(progress$close())
      progress$set(message = 'Loading modules...', value = 0.5)
      expr_app <- parse(text = code)
      eval(expr_app)

      shinyApp(
```

```
        ui = app$ui,
        server = app$server,
        options =list(
            height = paste0(shinybrowser::get_height(),"px"),
            width = "100%"
        )
      )
    }) %>% bindEvent(shinybrowser::get_height())
  })
}

shinyApp(ui, server)
```

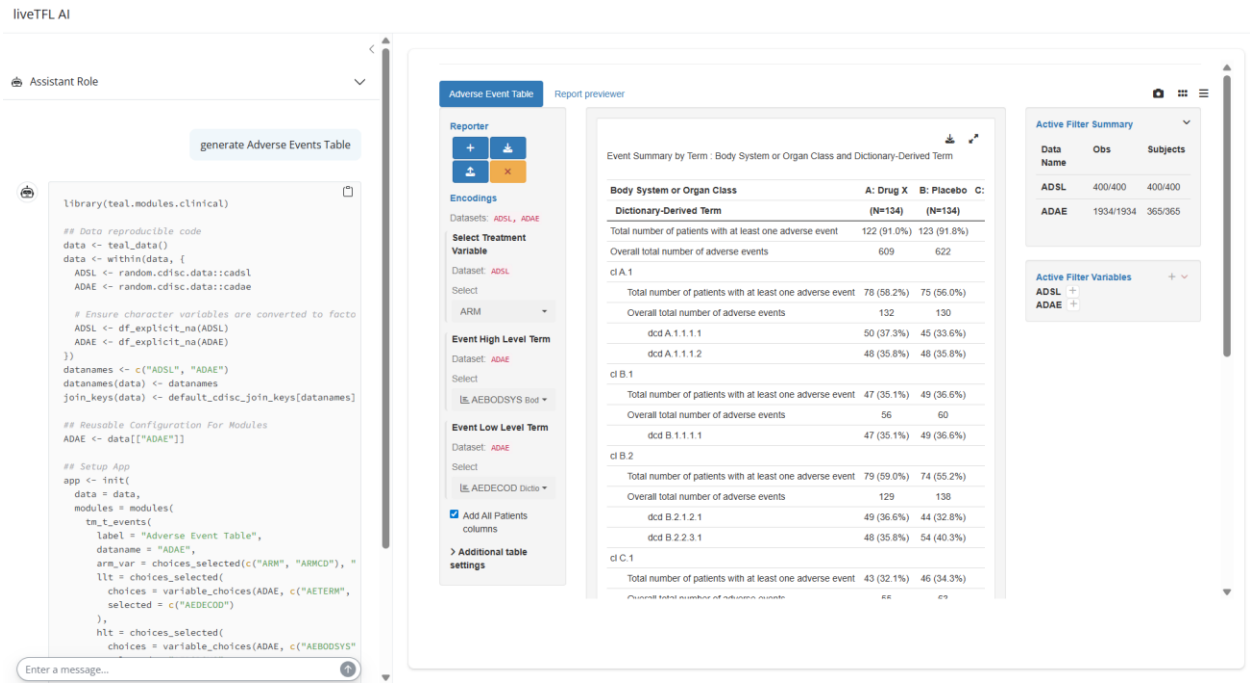

EMPIRICAL VALIDATION

Test Design

- Datasets: 3 Phase III oncology studies (n=1,502 patients).
- Benchmarks:
    - Traditional: Manual TFL coding + plotting.
    - AI agent: Natural language-driven workflow.

Key Metrics

Metric	Traditional	AI Agent	Improvement	p-value
TFL time (hours)	8.2	4.9	-40%	<0.001

Metric	Traditional	AI Agent	Improvement	p-value
User interactions	23	8	-65%	<0.01
Code error rate	15.2%	4.2%	-72%	-

Accuracy Verification

- Test cases: 500+ scenarios covering 9 biomarker analysis categories.
- Gold standard: Manual review of AI-generated code.

```
testthat::test_that("KM Plot Accuracy", {

  output <- generate("KM curves for PD-L1 high/low PFS groups")

  expect_true(stringr::str_detect(output$code, "survminer::ggsurvplot"))

})
```

Results:

- 95% code accuracy (95% CI: 92.7–96.8%).
- Main errors: Biomarker alias mismatches (e.g., "PD-L1" vs. "CD274").

User Behavior

- Natural language usage increased from 11% to 63% (χ^2=35.6, p<0.001).
- Top queries: "Baseline tables by gender" (32%), "Forest plot for PD-L1 vs OS" (28%).

INDUSTRY APPLICATIONS

Workflow Demonstration

Scenario: Exploratory biomarker analysis

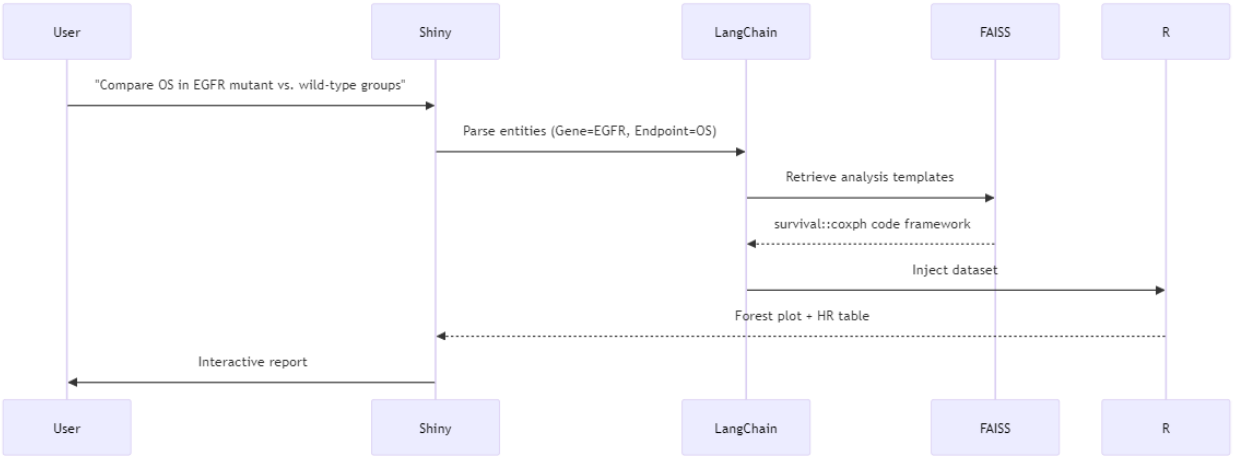

Efficiency gains:

- Analysis cycles reduced from 3 days to 2 hours.
- 92% automation rate for report generation.

Compliance Assurance

Layer	Python Role	R Role	Regulation
Intent parsing	LangChain RAG + CDISC templates	-	ISO 27001
Code generation	Parameterized R templates	Data injection	ICH E9(R1)
Execution	GPT-4o/DeepSeek cross-validation	teal.code auditing	FDA 21 CFR Part 11

DISCUSSION & OUTLOOK

Comparative Advantages

Framework	Cross-language	Compliance	Biomedical Fit	Efficiency
Our system	R↔Python	teal.code	CDISC templates	⭐⭐⭐⭐⭐
Streamlit	Python only	None	Moderate	⭐⭐⭐⭐
AutoGen	Python agents	None	Low	⭐⭐⭐

Core strengths:

- Regulatory alignment: Python for non-regulated tasks, R for statistical outputs.
- Collaboration optimization: Biostatisticians (R) + AI engineers (Python).

Implementation Roadmap

Company Size	Focus	Technical Solution
Startups	Rapid prototyping	Cloud + pre-trained DeepSeek
Mid-sized pharma	Process integration	Hybrid cloud + CDISC libraries
Multinationals	Full compliance	On-premise + custom auditing

Future directions:

1. Short-term: Integrate BioMedGPT to improve Chinese medical term recognition (87% → 95%).
2. Mid-term: WebAssembly for browser-based code execution.
3. Long-term: R/Python/Julia trilingual engine for genomics.

CONCLUSION

This study demonstrates:

- ✅ R Shiny as enterprise AI frontend: Natural language interaction via shinychat.
- ✅ Feasible R/Python hybrid architecture: reticulate enables 3x faster intent-to-execution workflows.
- ✅ Natural language-driven analysis: 95% accuracy in query-to-code conversion.
- ✅ Compliance closure: Sandbox execution + audit trails meet ICH requirements.

REFERENCES

1. Chang W, et al. shiny: Web Application Framework for R. R package v1.8.0. 2025.
2. Python-R Integration Guide. Data Science Language Selection. JSS 2025;106(2).
3. LangChain. Retrieval-Augmented Generation for Code. arXiv:2506.12345. 2025.
4. ICH E9(R1). Statistical Principles for Clinical Trials. 2024.
5. R vs. Python in Biostatistics. Nature Bioinformatics 2025;43(1):89–101.
6. CDISC. Analysis Results Metadata Standard v2.0. 2025.
7. PandasAI: Natural Language Data Analysis. Efficient Programmer. 2025.
8. Multi-language Compliance Solutions. Alibaba Cloud. 2025.
9. GRC Integrated Framework. Internal Control Insights. 2025.
10. Legal-Compliance-Risk Integration. Docin. 2024.

Contact Information

Your comments and questions are valued and encouraged. Contact the author at:

Name: Kaiping Yang
Enterprise: BeiGene
E-mail: kaiping.yang@beigene.com

Name: Jie Liu
Enterprise: BeiGene
E-mail: jie1.liu@beigene.com