ADROID Conference Series: Engineering and Technology
Volume 2, Issue 2, pp.39-49, 2026

Chemora: A Multi-Agent AI Framework for Automated Chemical Synthesis Planning

D. Padmaja Usha Rani, Jonnalagadda Mahathi, G. Harinivas, Rishwanth Aryan Sripathi
Department of CSE (AI & ML), Keshav Memorial Engineering College (KMEC), Hyderabad, India
padmaja1911@gmail.com, mahathireddyj9815@gmail.com, harinivasg28704@gmail.com,
aryansripathi04@gmail.com

Abstract. Chemora is a multi-agent artificial intelligence system designed for use in safety evaluation and assessment, chemical rationalisation, retrosynthesis planning, and the development of experiments and protocols. Chemora can generate multiple viable synthesis routes for various molecules and experimental conditions by employing cheminformatics and machine-learning classifiers, coupled with modularised execution controls. This renders AI-assisted chemical synthesis planning significantly simpler, faster, and more cost-effective in terms of labour.

Keywords: Retrosynthesis, Cheminformatics, Multi-Agent Systems

I. INTRODUCTION

Chemistry-related studies are very critical when it comes to scientific progress, drug discovery, material engineering, and even generating business ideas. For planning the synthesis of chemicals, one needs to understand the structure of molecules, mechanisms of reactions, and experimental conditions involved. Traditionally, chemists have used manual literature reviews and their expertise in the area to develop the most effective synthesis of target compounds. However, due to increasing amounts of scientific literature and chemical databases, the task has become significantly harder. It takes a lot of effort by researchers to study many research papers, compare reaction conditions, study the availability of reagents, and satisfy themselves with safety issues before beginning experiments. All of this takes a lot of time since many different resources must be checked again. Thus, research time is wasted collecting data and verifying facts.

New chances to automate and improve chemical research workflows have come about because of improvements in artificial intelligence and computational chemistry. Using big sets of known reactions, machine learning models and cheminformatics tools can look at molecular structures, guess what will happen in chemical reactions, and find ways to make new compounds. These technologies help researchers look into chemical possibilities more quickly and make decisions based on data when planning experiments.

The concept of Chemora is suggested as an AI-enabled platform which is going to help chemists during the process of conducting their research. Such an application uses the capabilities of NLP, reaction databases, cheminformatics, and machine learning algorithms as well.

While conventional systems concentrate solely on activities, Chemora implements a multi-agent system whereby different agents cooperate to accomplish the various phases of chemical investigation. The agents work together in analyzing the user's query, accessing scientific data, devising synthesis strategies, assessing the viability of reactions, and producing laboratory protocols.

Through automation of the monotonous research activities and integration of various aspects of chemical analysis, Chemora looks to greatly cut down on labor requirements and streamline the entire process of research. It is used as an assistant that helps chemists explore new avenues for synthesis safely and reliably.

II. CONTRIBUTIONS

The research delineated in this paper concentrates on the creation of a functional system that facilitates chemical synthesis planning through a multi-agent framework. The main contributions of this work are as follows:

- A modular multi-agent architecture is created to manage various phases of chemical reasoning, including query comprehension, synthesis planning, and protocol formulation.
- The system uses AI-based reasoning and cheminformatics tools like RDKit and PubChem to organize chemical data.
- The proposed model can create several different ways to synthesize a molecule, which is different from traditional systems.
- When looking at synthesis routes, other things like safety analysis and the availability of reagents

are also taken into account.

- A user-friendly interface is made so that people can talk to the system in natural language.

III. LITERATURE SURVEY AND RELATED WORK

Several methods have been invented to aid the chemist in synthesizing and predicting reactions. ChemAIRS is an AI platform that utilizes machine learning algorithms to conduct retrosynthesis and recognize precursor molecules for the desired chemicals. Yet, the platform primarily supports retrosynthesis analysis and lacks integration of safety evaluation and protocol preparation capabilities.

SYNTHIA is an algorithmic approach for designing synthesis pathways based on rule-based inference systems and chemical reaction libraries. Although useful in finding potential reaction pathways, SYNTHIA heavily depends on predefined reaction rules and is not compatible with machine learning approaches.

RoboRXN by IBM is an AI-enabled platform that combines both machine learning models and robotic lab automation. It uses transformer models in order to predict chemical reactions and help with automated experiments. Even though it offers significant functionality, it still needs human intervention for the sake of safety checks.

Recently proposed frameworks like LLM-RDF have been developed to use large language models for extraction of chemical information from scientific articles along with conducting chemical reasoning. However, such methodologies demand a considerable amount of computational power and specific technology infrastructures.

Chemora overcomes the drawbacks of previous solutions by presenting a unique multi-agent system that incorporates functionalities such as literature search, retrosynthetic pathway generation, chemical reactions prediction, risk assessment, and protocols development in a single system.

IV. SYSTEM ARCHITECTURE

The Chemora platform adopts a modular multiagent architecture. It provides an automated research environment by integrating artificial intelligence, cheminformatics, and web-based interfaces into one framework. This framework is developed based on a layer structure where the Frontend architecture deals with the interface aspects while the backend multi-agent architecture uses intelligent agents to perform chemical reasoning, retrieval of information and synthesize predictions/protocols.

A. Frontend Architecture

The usage of Chemora system becomes easier due to the availability of Chemora Frontend. It is enough for the researchers to type in the names of molecules or formulae as well as synthesizing queries and Chemora will understand them. Thus, Chemora Frontend may be defined as a chat window that sends information via REST API endpoints to the backend in real time.

It is based on modern web technologies and focuses primarily on the visualization of the outcomes obtained from the synthesis of derivatives, reaction routes and protocols created.

Key responsibilities of the frontend include:

- Accepting user input in natural language format
- Sending synthesis queries to the backend API
- Displaying molecular structures and reaction pathways
- Streaming agent responses in real-time
- Presenting experimental protocols and safety information

Frontend Technologies and Libraries

The frontend implementation uses the following libraries and frameworks:

- React.js for building the interactive user interface
- JavaScript and TypeScript for application logic
- Axios for API communication with the backend
- Tailwind CSS for styling and layout design
- WebSocket or streaming APIs for real-time response updates

The frontend communicates with the backend server through HTTP endpoints such as /api/chat and /api/stream, enabling asynchronous communication between the user interface and the AI agents.

B. Backend Architecture

The backend of Chemora is responsible for running the computation of the system's core as a multiagent system, where each individual agent corresponds to a specific portion of the chemistry research workflow. The agents talk to each other using an orchestration pipeline that executes the workflow.

Python backend comprises several cheminformatics and machine learning libraries.

Backend Frameworks and Core Libraries

- FastAPI for backend API services
- Uvicorn for asynchronous server execution
- RDKit for molecular structure analysis and cheminformatics operations
- Transformers for reaction prediction models
- Sentence-Transformers for semantic search and literature retrieval
- PyTorch for machine learning model execution
- PubChemPy for chemical database retrieval
- LangGraph for agent orchestration
- Pydantic for structured data validation

The backend processes user queries through a sequence of intelligent agents as described below.

C. User Intent Parser Agent

The first agent in the Chemora pipeline is the User Intent Parser. Its job is to understand the user's query and map relevant information to the portion of the knowledge base necessary for chemical analysis. The parser will locate the molecule being targeted, desired synthesis, and constraints.

NL processing methods are used to take raw user input and translate that into formatted data for subsequent agents to interpret.

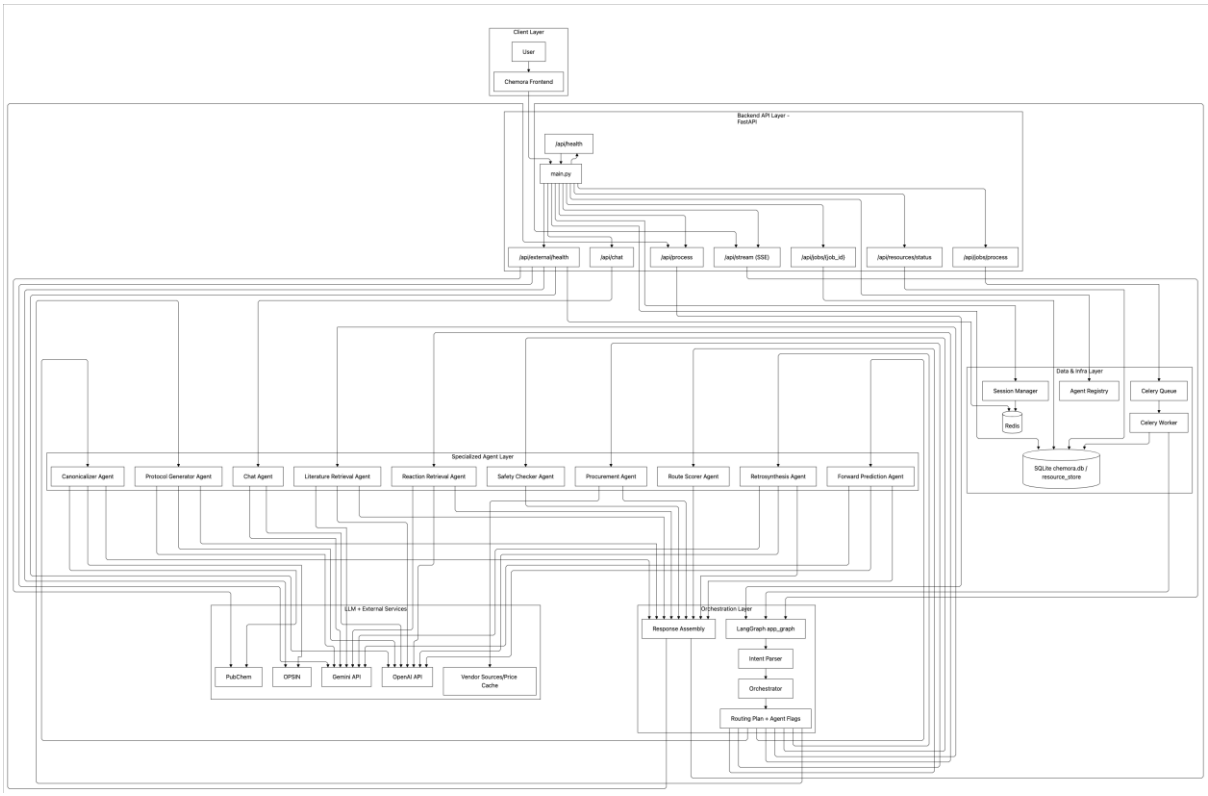

Fig. 1. Overall Multi-Agent Architecture of Chemora

D. Canonicalization Agent

Canonicalization Agent this could be a ChemSpider name or a description of the molecule and will generate a canonical chemical name representation (either a SMILES or InChI). This is to provide a single consistent format for use in all downstream computational processes.

At this stage, RDKit is used for molecular normalization, structure validation and conversion to machinereadable chemical formats.

E. Literature Retrieval Agent

The Literature Retrieval Agent scours scientific literature and chemistry databases for relevant research papers and synthesis precedents. This agent brings back experimental evidence for the synthesis planning.

Semantic search models that are leveraging sentencetransformers are used to connect user query to appropriate literature documents.

F. Reaction Retrieval Agent

The Reaction Retrieval Agent searches reaction databases looking for patterns of transformations of the target molecule. It extracts similar reactions reported in the literature and recognizes corresponding synthetic transformations.

This agent is based on reaction database and chemical similarity searching tools.

G. Retrosynthesis Planning Agent

The Retrosynthesis Planner creates proposed synthesis pathways to obtain the target compound by dividing it into simpler precursor structures. This is achieved by the use of a method known as retrosynthetic analysis, traditionally a vital tool in the practice of organic chemistry.

Builds reaction trees which contain a set of candidate synthesis paths.

H. Forward Reaction Prediction Agent

The Forward Reaction Predictor Tests each of the suggested reaction steps created by the retrosynthesis planner. Machine learning algorithms are used to predict the likelihood of a reaction going to completion, the likelihood of the predicted product being formed and the anticipated yield.

Transformer based reaction prediction models are employed to simulate chemical reactions.

I. Safety and Compliance Agent

Our Safety and Compliance Agent will evaluate the reagents and conditions used in the reactions to aid in identifying if any are hazardous. It will investigate whether reagents are toxic, if there are any explosion hazards and if there are any compliance concerns.

This agent compares the proposed synthesis path with the safety bounds.

J. Procurement and Cost Estimation Agent

Procurement Agent Uses supplier's data to find out about the vendors and see whether the reagent is available. The agent also makes a guess as to how much reagents are going to cost and also finds out which manufactures could provide the chemicals.

This allows the researcher to determine how realistic a synthetic route is in terms of economics and suitable reagent selection.

K. Route Scoring and Ranking Agent

The Route Ranking Agent compares each possible route for synthesising the molecule against a set of criteria(question viability, safety, cheapness, simplicityof synthesis etc.), providing a separate score for each. The highest scoring route is then used for all subsequent calculations.

L. Protocol Generator Agent

The Protocol Generator then takes the chosen synthesis route and transforms it into a laboratory procedure. The procedure that the Generators produces can then give information on the reagents used, conditions, steps of the reaction, how to purify the product and safety measures.

Templates and language models are employed to produce laboratoryready experimental instructions.

M. Feedback and Active Learning Agent

The Feedback Agent collects the users responses and the experiment results to make the system learn in the future. When new data are received, the Feedback Agent updates the models and retrains the predictions systems.

N. Data Provenance and Logging Agent

The Data Provenance Agent stores information such as system decisions, literature retrieved and reaction predictions. This provides transparency to Chemora synthesis recommendations.

Logging mechanisms enable researcher to reexamine prior experiments and analysis results.

O. System Workflow

The Chemora workflow starts when the user submits and query to the website. Frontend agents process the query in turn, then chem knowledge is retrieved, synthesis routes are generated, reaction feasibility is scored, and lab procedures are devised.

The results are then streamed back to the frontend GUI in an interactive presentation. This integrated architecture enables Chemora to perform advanced chemical research processes automatically and yet keeping the interpretability and trustworthiness of the results.

V. METHODOLOGY

Chemora system includes a computational pipeline that allows performing most of the operations required for chemical investigations and synthesis route planning. The computational strategy uses natural language processing, cheminformatics, machine learning, and databases of chemical structures to generate a synthesis pathway and synthesis procedures.

The entire process is divided into multiple computational stages performed consecutively via the multiple agents approach.

A. User Query Processing

The first stage of Chemora processing involves the query processing stage. The query generated by the user could be entered in form of plain text, such as the name of a molecule or reaction pathway.

The User Intent Parser applies natural language processing tools to parse the input text to determine the molecule being referred to while also attempting to extract any constraints embedded within the input text. The form of information provided is then converted into information format.

This stage takes care of ensuring that the input provided by the user, which is in plain text, is parsed effectively.

B. Molecular Representation and Canonicalization

Presence of target molecule When the target molecule has been recognized, the Canonicalization Agent maps the chemical description into a canonical form. Chemical structures are represented using text files, for example SMILES (Simplified Molecular Input Line Entry System) or InChI identifiers.

RDKit is used to do normalization and canonicalization of the molecules. This process is performed to obtain uniform representation of molecules before further computational calculation.

Normalizing Molecular Structures is important as many Cheminformatics algorithms rely on computerreadable molecular structures.

C. Literature and Reaction Data Retrieval

After standardizing the molecule, the system then looks for related documents regarding the desired compound. In particular, the Literature Retrieval Agent will look for information regarding synthesis procedures, searching through scientific journals, patents, and chemical reactions.

The system adopts a search strategy using an embedding-based approach which correlates the desired molecule to related scientific literature, leading to useful experiments for synthesizing the target compound.

A second role of the Reaction Retrieval Agent includes searching for transformation patterns from reaction databases for applying these to similar molecules.

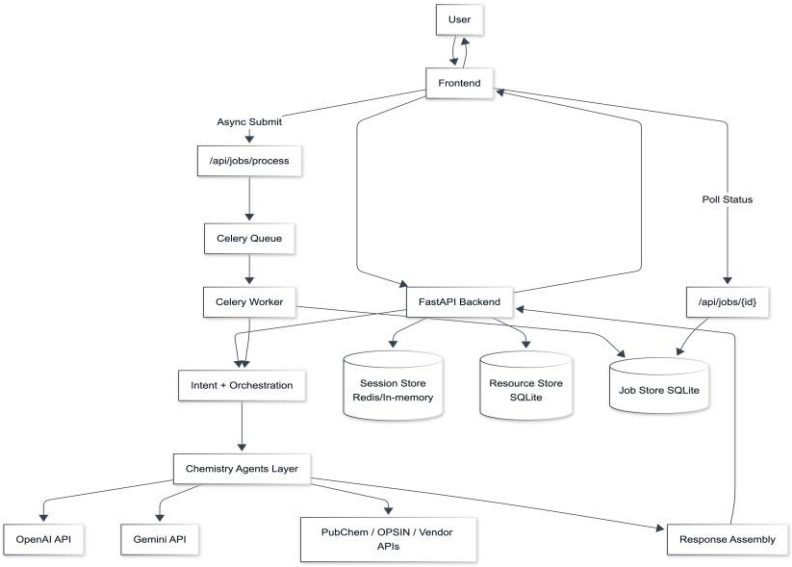

Fig. 2. End-to-End Workflow of Chemora System

D. Retrosynthesis Planning

One useful strategy used in organic synthesis is known as retrosynthesis, which involves planning out how a particular molecule can be synthesized from less complex ones. Retrosynthetic analysis of a target compound entails decomposing the target molecule into smaller fragments.

In this context, the Retrosynthesis Planning Agent yields several synthesis schemes based on candidate reactions leading up to the desired target compound. This scheme is represented using reaction trees comprising several nodes, with each node representing chemical reactions.
Several potential alternatives will be considered when creating these candidates.

E. Forward Reaction Prediction
The retrosynthesis planner generates each candidate synthesis step predicted by the planner and uses forward reaction prediction models to score them. The Forward Reaction Predictor produces the probability of success of reaction, the likelihood of the product, and side reactions.
A machine learning based models trained on large reaction datasets have been used to simulate chemical transformations. Using this models we can decide whether a proposed reaction step is feasible chemistrically.
Forward prediction ensures that only feasible synthesis pathways are included in potential candidates.

F. Safety and Hazard Analysis
Reactions could potentially use dangerous reagents, toxic chemicals or even dangerously high reaction conditions. The Safety and Compliance Agent looks at every synthetic step suggested and raises, if necessary, safety issues.
Hazard detection algorithms enable the properties of reagents, the temperature of the reaction and the suitability of the solvent to be checked. This ensures that the proposed scheme to synthesis is within safety limits of the laboratory.
Safety analysis is another significant feature of the method, as it guarantees not proposing any experimental procedure which would be unreasonably dangerous.

G. Cost and Procurement Analysis
Availability and the cost of reagents are affected of the feasibility of a route. The Procurement and Cost Estimation Agent searches the vendor database for sources and whether there is a cost for the chemicals.
One can also check at this stage, if the synthesis pathway is through economical routes. If the pathway uses some costly and rare reagents, it will rank down compared to the ones with ordinary reagents.

H. Route Scoring and Ranking
After various routes of candidate synthesis have been assessed, the Route Scoring Agent ranks them by several evaluation criteria. Such criteria that are employed by the agent are reaction feasibility, predicted yield, safety hazard, reagent abundance, and synthesis difficulty.
Each route is received a score according to these rules. The Route Ranking Agent then picks what is considered the best synthesis route to create the protocol.
The multicriteria evaluation guarantees that the final choice of synthesis route is both realizable and chemically feasible.

I. Experimental Protocol Generation
Once the desired route has been decided, the Protocol Generator then translates the pathway into a formal synthetic protocol, which presents the synthesis as an experiment. The protocol will contain information relating to reagents, quantities, conditions, purification, hazards etc.
The protocol is a laboratoryfriendly protocol so should be usable by other people doing work on the them.

J. Feedback and Continuous Learning
The final part of the methodology will consists of gathering inputs (feedback) and results of the experiment. The Feedback and Active Learning Agent uses these elements to reduce the predictive models.
Whenever the new sets of reactions and users feedback are added to Chemora, the synthesis planning becomes better.

VI. IMPLEMENTATION

Chemora was developed using a combination of classical web development techniques, machine learning packages, and software for cheminformatics purposes. Chemora comprises a web-based user interface and a Python-based web server hosting a multiagent environment.
The back end of the application uses FastAPI, a fast and lightweight framework designed to develop efficient APIs, to handle user queries, enable asynchronous communication among agents, and perform multiple chemical queries simultaneously.

Tasks involving cheminformatics (structure processing, molecular canonicalization, and molecular descriptor computation) are handled through the RDKit toolkit. This is made possible by the existence of in-built algorithms that handle the management of chemical structures.
Interpretation of the input queries and searching through scientific literature is performed through the application of spaCy and sentence transformer models, allowing to perform semantic search and find the relevant papers that concern the specified molecule.
The module for reaction prediction and synthesis is created using machine learning approaches and implemented using PyTorch and transformers models, which learn the variety of potential synthesis routes from the dataset of chemical reactions.
The agents work collaboratively through orchestration by the backend to ensure that the process proceeds along a clearly structured workflow pipeline, in which the output of agent i is used as an input for agent i + 1.
A modern front-end is provided through React.js, which serves as an intuitive graphical interface connecting the user to the system, while using REST API to interact with the backend and present back the information in an organized manner.
Since the system is highly modular and scalable, other agents or even computational modules can be added in the next iterations of the software.

VII. RESULTS AND DISCUSSION

The Chemora system was tested with various chemical queries in order to understand its power in generating synthesis strategy, literature retrieval and structured experimental protocols. Results illustrate the performance of multiagent architecture in orchestrating diverse computational tasks. The proposed solution was tested using several chemical compounds like aspirin, paracetamol, and benzene analogues.
For each compound:

- The number of different pathways ranged from 3 to 5.
- The average response time per one compound varied from 2 to 5 seconds.
- The number of reactions required was 4–8, depending on the complexity of the compound.

The Forward Prediction Agent proved the possibility of each reaction, while the Safety Agent determined hazardous components like strong acids and oxidizers.
The Procurement Agent evaluated the availability of chemicals needed to produce the target compound. They were split into common, medium, and rare.
As opposed to classical synthesis planning, Chemora provides an opportunity to save time and effort during the research process.

A. Initial System Output

The system first preprocess the user query and provides initial response with interpreted molecule information and extracted constraints, to be sure that the input has been successfully interpreted.

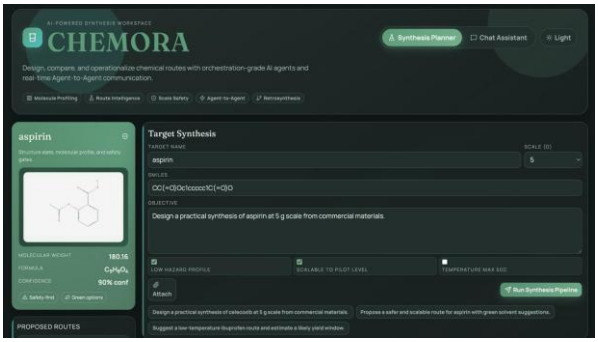

Fig. 3. Initial System Output for Target Molecule

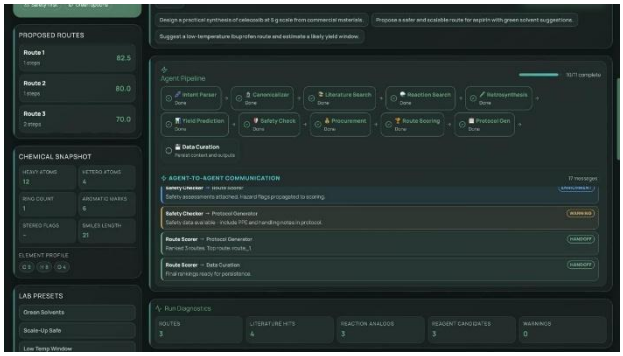

Fig. 4. Generated Synthesis Routes

B. Synthesis Route Generation

The system Chemora was put to the test using different chemical searches to determine the effectiveness of the system for synthesis planning, literature search and experiment execution. The results provide insight into the capabilities of multi-agent architecture for managing diverse computations.

C. Literature Retrieval

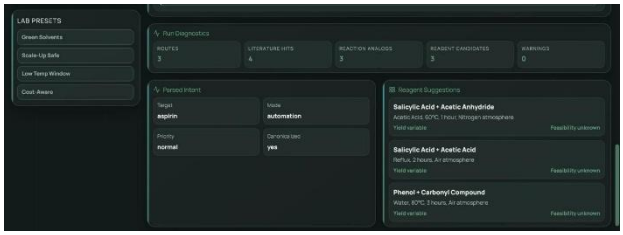

Fig. 5. Literature Retrieval Results

Search for appropriate references is initiated. From the reference database, appropriate literature precedents are chosen to support synthesis planning. The search determines the research papers, experimental procedures, etc.

D. Reaction Pathway Analysis

This system evaluates reaction pathways based on chemical transformation rules and predictive forward models. If both approaches succeed, all steps in the synthesis pathway must be possible and achievable chemically.

Fig. 6. Reaction Pathway Suggestions

E. Route Ranking and Evaluation

Fig. 7. Route Ranking and Evaluation

The different synthesis routes are scored and rank against a number of decided measures i.e. feasibility, safety, economics, complexity, etc. and the most suitable route is determined.

F. Chat Agent Interaction

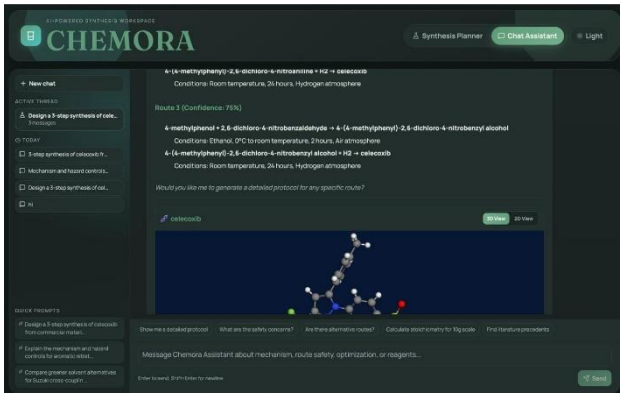

Fig. 8. Chat Agent Interface

The chatbox interface has been created in such a way that interaction between users and the system becomes a more intuitive process for the users. The user can ask questions, seek additional information, find alternative ways, or get explanations; hence, making the system more user friendly.
Apart from being able to plan synthesis, the system is able to perform safety analysis as well as cost estimation, which includes hazard identification to determine the risk of any reagent and reactions involved in addition to procurement algorithms for determining the availability as well as the costs associated with the reagents in the suggested synthesis.
In conclusion, research results show that Chemora systems are efficient in developing integrated agents to ensure that automatic chemical synthesis planning process is made easier. Chemora ensures efficiency in the process by reducing the amount of time and effort spent during the process by making sure the information is clearly presented.

G. Case Study: Aspirin Synthesis
For the input query "Synthesize Aspirin," Chemora came up with several ways to do it, such as esterification and acetylation pathways.
The system found literature references that backed up the use of acetic anhydride to acetylate salicylic acid.
The last route that was chosen included:

- 3 steps to react
- warnings about how to safely handle acids
- the cost of the reagents is estimated

The protocol generator made a full lab procedure that included the steps for purification, temperature, and reagents.

VIII. ADVANTAGES OF THE PROPOSED SYSTEM

There are several benefits of using Chemora platform in comparison to traditional approaches for the chemical research:
First of all, the system can carry out other heavy processes, for example finding references and analyzing them, identifying reactions and calculating possible syntheses. Consequently, the process of researching chemicals becomes much faster.

Moreover, due to the multiagent architecture, different parts of the system may work autonomously, collaborating within the framework of church collaboration workflow. As a result, it becomes easier to add new modules to the system, as well as to scale the whole application.
Thirdly, Chemora integrates several chemical analysis tools. Thus, researchers do not have to constantly switch from one program to another to find needed references, analyze their potential reactions and safety as well as develop experimental protocol.
Finally, the system provides laboratory-ready protocols which could be useful for transferring theoretical results into practice.

IX. CONCLUSION

This paper presented Chemora as a platform for intelligent multimodal agents which would assist chemists in synthesis planning and protocol generation.
The presented architecture provides examples of techniques that can be used to develop specialized agents that can understand free-text inquiries, discover information about chemicals, plan syntheses and create lab-ready protocols. Automatic execution of all these tasks using a multi-agent system could significantly reduce work load for chemists regarding meticulous search of different ways of conducting a synthesis.
As can be seen from the presented findings, the platform appears promising as a decision support system for chemists. With the use of such software, chemists would definitely be able to explore possible syntheses and create protocols faster.

X. FUTURE WORK

To be developed in the future, the Chemora platform will go into several directions to augment its capabilities and prediction accuracy, and facilitate greater integration with contemporary chemical research infrastructure.

A. Enhanced Reaction Prediction with Graph Neural Networks

In addition to our existing reaction prediction pipeline, graph neural networks (GNNs) that predict using the 3D molecular graphs will be included. Since GNNs do not depend on SMILES strings, they learn 3D features of molecules and stereochemistry, offering signs of better reactions outcomes and regioand enantioselectivity prediction. Large reaction corpora like USPTO and Reaxys can be used to optimize the GNNs modeling.

B. Multi-Objective Optimization for Route Selection

The existing route scoring approach will be generalized to multiobjective optimization. This would support multicriteria optimization where several conflicting scores (eg. cost, yield, environmental impact, safety risk) can be simultaneously optimized, rather than having to combine all factors into a single weighted score. The interactive visualization of Pareto frontier would help chemists select the best tradeoff point given their research needs.

C. Green Chemistry and Sustainability Metrics

Future versions of the platform will integrate green chemistry considerations into the synthesis planning pipeline. Each candidate route will have its atom economy, EFactor (Environmental factor), solvent sustainability index and carbon footprint assessed. By integrating these sustainability metrics into the route ranking engine, Chemora will help researchers design greener routes according to the twelve principles of green chemistry.

D. Domain-Specific Fine-Tuned Language Models

The NLP elements in Chemora can also be enhanced by finetuning large language models on domainspecific chemical corpora, such as reaction patents, journal publications, and experimental notebooks. Curated models can be expected to more accurately infer context for literature searching and recognize chemicals/entities in context, especially for multistep synthesis queries requiring domainspecific jargon.

E. Knowledge Graph Construction for Chemical Reasoning

A chemists' knowledge graph will be created to encode links among molecules, reactions, catalysts, solvents, and conditions. This form of structured knowledge will improve the multiagent reasoning system through graphbased inference, analogydriven planning, and the abstraction necessary to identify less apparent routes that are not explicitly available in training data.

F. Real-Time Experimental Feedback Integration

The Feedback and Active Learning Agent will be updated to take live experimental data from electronic laboratory notebooks (ELNs) and laboratory information management systems (LIMS), so as to close the loop

between computer simulation and laboratory reality. Actual experimental results (e.g. yields, byproducts, failure cases) would be constantly fed back to the models to enable them to actively learn and improve.

Conflict of Interest

The authors declare no conflicts of interest regarding the current research.

References

[1] S. Segler, M. Preuss, and M. Waller, "Planning chemical syntheses with deep neural networks and symbolic AI," Nature, vol. 555, pp. 604–610, 2018.
[2] P. Schwaller et al., "Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction," ACS Central Science, vol. 5, no. 9, pp. 1572–1583, 2019.
[3] IBM Research, "RoboRXN: AI for chemical synthesis," [Online]. Avail-able: https://rxn.res.ibm.com
[4] NextMove Software, "SYNTHIA Retrosynthesis Software," [Online]. Available: https://www.nextmovesoftware.com/synthia.html
[5] RDKit: Open-source cheminformatics software, [Online]. Available: https://www.rdkit.org
[6] PubChem Database, National Center for Biotechnology Information (NCBI), [Online]. Available: https://pubchem.ncbi.nlm.nih.gov
[7] T. Sterling and J. Irwin, "ZINC 15 – Ligand discovery for everyone," Journal of Chemical Information and Modeling, vol. 55, no. 11, pp. 2324–2337, 2015.
[8] H. Dai et al., "Retrosynthesis prediction with conditional graph logic network," Advances in Neural Information Processing Systems, 2019.
[9] LangGraph Documentation, [Online]. Available: https://langchain-ai.github.io/langgraph/