# Planning Ontology: A Tutorial on Knowledge Representation and Explainable Planning

**Bharath Muppasani**[1], **Nitin Gupta**[1], **Biplav Srivastava**[1], **Vignesh Narayanan**[1],

[1]Artificial Intelligence Institute at University of South Carolina,
bharath@email.sc.edu, niting@email.sc.edu, biplav.s@sc.edu, vignar@sc.edu,

### Abstract

Ontologies provide a rich, semantically grounded framework to represent complex domains. The Planning Ontology models fundamental planning concepts—domains, problems, planners, plans, and action explanations—as OWL classes and properties. We demonstrate how to leverage this knowledge graph for three key applications: (i) SPARQL-driven planner selection based on domain features and past performance, and (ii) extraction and composition of human-readable explanations for plan actions. Through detailed code examples and guided Jupyter notebooks, participants will integrate ontology queries into Python-based planning pipelines and evaluate query efficiency. By the end of the session, attendees will be equipped to use ontological insights to improve planner choice, transparency, and reproducibility in AI planning workflows.

**Keywords:** Automated Planning, Neuro-symbolic Approaches, Ontologies, Knowledge Graphs.

## Introduction

Automated planning, a core sub-field of Artificial Intelligence (AI), focuses on finding a sequence of actions that transition an agent from an initial state to a desired goal state (Ghallab, Nau, and Traverso 2004). The ability to generate plans for complex domains has led to significant progress in real-world applications, from optimizing logistics and manufacturing (Behandish, Nelaturi, and de Kleer 2018) to enabling sophisticated robotics (Karpas and Magazzeni 2020). The field is characterized by a rich and ever-expanding ecosystem of planners, search algorithms, heuristics, and diverse planning domains. This wealth of options, however, presents a significant bottleneck.

Traditionally, improving planner performance or selecting the right planner for a new problem relies on manually exploring and curating combinations of these components—a process that is time-consuming and lacks scalability (Muppasani et al. 2024). The core issue is the "absence of a structured framework" to represent the critical meta-information about planner capabilities and domain features, which limits our ability to reason efficiently about their suitability for specific tasks.



To address this gap, this tutorial introduces a hands-on approach centered around the **AI Planning Ontology**. Ontologies provide a formal representation of concepts and their relationships, enabling systematic reasoning, knowledge sharing, and reuse (Guarino, Oberle, and Staab 2009). While previous ontologies for planning like PLANET (Gil and Blythe 2000) have been proposed, they have had limited reusability and scope. The AI Planning Ontology, which this tutorial is built upon, offers a more comprehensive, open, and extensible model designed to support modern planning challenges.

In this tutorial, we focus on practical applications built atop a pre-constructed knowledge graph derived from this ontology. Our objectives are to illustrate how semantic queries can be used to guide two key applications: (i) automated, data-driven planner selection and (ii) the generation of human-readable explanations for plans. We will motivate each use case with competency questions, detail the SPARQL patterns needed to query the knowledge graph, and provide hands-on experience integrating these queries into Python-based planning pipelines. By the end of this session, attendees will be equipped to use ontological insights to improve planner choice, transparency, and reproducibility in their own AI planning workflows.

## Goal of the Tutorial

**Target Audience:** This tutorial is designed for graduate students, researchers, and AI practitioners who work with automated planning systems and are interested in leveraging knowledge representation to overcome practical challenges in planner selection and explainability. It is ideal for attendees who know basic Python and want to learn how semantic technologies can be applied to symbolic reasoning tasks.

**Takeaways for the Audience:** Participants will leave with a working Python pipeline and the skills to:

- Write SPARQL queries to retrieve and rank planners based on domain requirements and past performance.
- Integrate semantic query results into a planner selection and execution loop.
- Construct coherent, human-readable plan explanations by composing query results about plan steps and action semantics.
- Understand the value of formally modeling planning artifacts to enable new reasoning capabilities.

Table 1: Detailed Outline of the Tutorial

| Time | Activity | Presenter |
|---|---|---|
| 10 mins | **Introduction to the Tutorial:** Objectives, materials overview, environment setup. | Biplav Srivastava |
| 25 mins | **Deep Dive into the Planning Ontology & AI Planning:** OWL schema review, key classes/properties, competency questions. | Bharath Muppasani |
| 20 mins | **Use Case 1: SPARQL-Driven Planner Selection:** Introduction and significance; query design and optimization; notebook exercise integrating planner choice into Python. | Bharath Muppasani |
| 25 mins | **Use Case 2: Plan Explanation Extraction:** Introduction and significance; SPARQL for action steps and explanations; notebook exercise on narrative assembly. | Nitin Gupta |
| 15 mins | **Discussion:** Survey related-work, Tooling, ma extension; discussion of resource extensions. | Biplav Srivastava |
| 5 mins | **Q&A Session:** Addressing queries and clarifications from participants. | All presenters |

## Innovation Aspect

The tutorial's core innovation is its knowledge-driven approach to solving planning tasks. While planner performance data and domain models are widely available, they exist as static, disconnected artifacts (e.g., IPC reports, PDDL files). Our approach is novel because it transforms this latent knowledge into an explicit, queryable RDF knowledge graph. This enables a dynamic, "ask-me-anything" paradigm for planning workflows, where decisions about which planner to use or how to explain a plan are guided by integrated, machine-readable data rather than by manual inspection and heuristics.

## Core Concepts, Methods, and Modeling Frameworks

The tutorial is built on a framework that combines ontological modeling with query-driven application logic.

- **Ontological Modeling (OWL):** We formally model concepts like `PlanningDomain`, `Planner`, and `Plan`. We use object properties to define relationships (e.g., a `Planner solvesRequirement` a domain `Requirement`) and data properties to store metadata (e.g., a `Plan hasPlanCost` a specific integer value).
- **Knowledge Graph Population:** We demonstrate how the ontology is populated with instances from public datasets like the IPC, transforming benchmark results into queryable semantic triples (e.g., `domain:blocksworld hasHighRelevance planner:FastDownward`).
- **SPARQL-driven Reasoning:** We use SPARQL as a reasoning mechanism to query the knowledge graph. The hands-on tutorial focus on specific query patterns for:
  - **Selection:** Filtering planners based on required domain features (e.g., ':typing', ':adl').
  - **Ranking:** Ordering planners using performance data (e.g., 'hasHighRelevance', 'hasMediumRelevance').
- **Explanation Framework:** We use the ontology's structure to support explainability. This involves querying for the sequence of `Step` instances in a `Plan` and retrieving their associated `hasActionExplanation` property to build a narrative.

## Illustrative Scenarios

The hands-on tutorial use two illustrative scenarios to make the concepts concrete.

- **Planner Selection:** Attendees are tasked with selecting a planner for classic AI planning problems like 'blocksworld'. They will write a SPARQL query to find all planners in the KG that support the ':typing' requirement and then rank them based on their historical relevance score for that domain. This demonstrates how to automate a choice that is typically based on developer experience or trial-and-error.
- **Plan Explanation Generation:** Given a pre-generated plan, attendees will write SPARQL queries to first retrieve the ordered sequence of actions. They will then perform a second query to fetch the human-readable explanation for each of those actions and compose them into a clear, step-by-step narrative of the solution.

## Technical Setup and Tools

To ensure a seamless and productive hands-on experience, this tutorial will be conducted using Google Colab, eliminating the need for any local software installation.

- **Participant Requirements:** Participants will need a laptop with a modern web browser and a Google account to access the Google Colab environment. No local installation of Python, libraries, or other software is required.
- **Materials Provided by Organizers:** We will provide a series of Google Colab notebooks that come preconfigured with all necessary Python libraries. Planning Ontology (Muppasani et al. 2024) and the tutorial's Knowledge Graph (KG), populated with multiple IPC instances, will be hosted and made accessible for querying directly within the Colab environment.

## Detailed Outline of the Lab

The lab is carefully designed for a duration of **1 hour and 45 minutes** to ensure a comprehensive yet focused learning experience. In this hands-on tutorial, we begin with an *Introduction to the Tutorial* where participants are oriented to the objectives, materials (Google Colab notebooks and the hosted Knowledge Graph), and the overall flow of the session. We then proceed to a *Deep Dive into the Planning Ontology & AI Planning*, reviewing the OWL schema, key classes and properties (e.g., `Planner`, `Plan`, `Action`), and the competency questions.

Next, in *Use Case 1 – SPARQL-Driven Planner Selection*, we motivate why automated planner choice matters, demonstrate how to write and optimize a SPARQL query to rank planners by feature support and historical performance, and walk through a live Google Colab notebook exercise integrating the query into a Python planning loop. In *Use Case 2 – Plan Explanation Extraction*, we show how to retrieve ordered action steps and their human-readable explanations from the ontology, and assemble them into a coherent narrative via Python code.

We conclude with a *Discussion and Q&A Session*, surveying related work and brainstorming ways to extend the tutorial pipeline—adding new domains, richer explanations, or dashboard visualizations, and addressing queries from participants.

## Participants and Prerequisites

**Estimated Number of Participants:**
We anticipate an audience of approximately 15 to 30 participants. This size is ideal for a hands-on lab, fostering an interactive environment where personalized guidance and interaction are possible.

**Prerequisite Knowledge:**
No specific prior knowledge is mandated, and we welcome all individuals with an interest in the topic. However, a basic understanding of AI planning concepts (e.g., domain, problem, action, plan) and introductory-level familiarity with Python programming will enhance the learning experience. The tutorial is designed to be self-contained, and our primary aim is to create an inclusive environment where all participants can gain valuable insights and practical skills.

## Supplementary Materials

1. **GitHub Repository:** All tutorial materials, including the ontology, scripts, and notebooks, are publicly available.
   - https://github.com/ai4society/planning-ontology
2. **Ontology PURL and Documentation:** The ontology is accessible via a persistent URL and includes comprehensive documentation (Ai4s 2023).
   - PURL: https://purl.org/ai4s/ontology/planning
3. **Hands-on Google Colab Notebooks:** The tutorial will be conducted through a series of interactive notebooks, each designed to answer specific competency questions using the provided Knowledge Graph (KG).
   - **01-Planner-Selection.ipynb:** Focuses on using the ontology for intelligent planner selection.
     - *Example Question:* ”Which planner has the highest relevance score for the ‘blocksworld‘ domain?”
   - **02-Plan-Explanation.ipynb:** Demonstrates how to extract information for generating human-readable plan explanations.
     - *Example Question:* ”For a given plan, what are its sequential action steps and their human-readable explanations?”

## References


Ai4s. 2023. Planning Ontology. https://github.com/ai4society/planning-ontology/. Accessed: 2024-09-26.

Behandish, M.; Nelaturi, S.; and de Kleer, J. 2018. Automated process planning for hybrid manufacturing. *Computer-Aided Design*, 102: 115–127.

Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*. San Francisco, CA: Elsevier. ISBN 978-1-55860-856-6.

Gil, Y.; and Blythe, J. 2000. PLANET: A shareable and reusable ontology for representing plans. In *AAAI Workshop on Representational Issues for Real-world Planning Systems. Website: https://www.isi.edu/b̃lythe/planet/*, volume 114.

Guarino, N.; Oberle, D.; and Staab, S. 2009. What is an ontology? *Handbook on ontologies*, 1–17.

Karpas, E.; and Magazzeni, D. 2020. Automated planning for robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1): 417–439.

Muppasani, B.; Gupta, N.; Pallagani, V.; Srivastava, B.; Mutharaju, R.; Huhns, M. N.; and Narayanan, V. 2024. Building a Plan Ontology to Represent and Exploit Planning Knowledge and Its Applications. In *Eighth International Conference on Data Science and Management of Data (CODS-COMAD’24), India*.

# Bharath Muppasani

🏠 510 Heyward St, Apt 108, Columbia, South Carolina - 29201

📫 bharath@email.sc.edu LinkedIn: bharathmuppasani 🔗 Website: bharathmuppasani.github.io


## Education

**University of South Carolina** — **Columbia, SC, USA**
*Doctor of Philosophy, Computer Science* — Jan 2022–Present
Advisors: Dr. Biplav Srivastava & Dr. Vignesh Narayanan

**International Institute of Information Technology** — **Naya Raipur, Chhattisgarh, India**
*Bachelor in Technology, Electronics and Communication* — Aug 2016 - May 2020

## Research Interests and Skills

– **Research Interests:** Reasoning (Automated Planning), Learning (Deep Learning, Reinforcement Learning, Large Language Models), Representation (Ontology).
– **Proficient:** Python, PDDL, C, SQL, SPARQL, Machine Learning Frameworks

## Selected Projects

**Planning Strategies for Dynamic Opinion Networks** — May 2023–May 2024
*Related Publications:* NeurIPS-24, AAAI-24 Demo Track
– Designed intervention strategies using ranking algorithms and neural network classifiers to identify key nodes for accurate information dissemination in dynamic networks.
– Developed a reinforcement learning framework for large-scale networks, analyzing five reward structures tailored to network dynamics.
– Conceptualized information spread as a numerical planning problem and built a network visualization tool for simulating information propagation and visualizing plans for targeted information spread.

**Solving Automated Planning Problems using Large Language Models** — May 2022–July 2023
*Related Publications:* ICAPS-24; IJCAI-23 Demo Track; NeurIPS-23 GenPlan
– Spearheaded the development and curation of a comprehensive dataset for fine-tuning and evaluating the performance of various Large Language Model (LLM) architectures, including encoder-only, decoder-only, and encoder-decoder configurations in automated planning scenarios using Planning Domain Definition Language (PDDL).
– Played a pivotal role in executing evaluations of LLM outputs, providing critical insights that informed the design and optimization of learning-based planners, notably the Plansformer model, as evidenced by its demonstrable efficacy accessible at Plansformer Website

**A Planning Ontology to Represent and Exploit Planning Knowledge** — May 2022–May 2024
*Related Publications:* ICAPS-23 PLATO; CODS COMAD-24
– Developing a planning ontology to leverage knowledge from diverse planning domains and improve planner performance, with a focus on action ordering constraints and open access resources.
– The documentation for the current stage of Planning Ontology is available on PyLODE and the OWL file is accessible via PURL.

**Solving complex puzzles with automated planners using symbolic representations** — Jan 2022–Dec 2022
*Related Publications:* ICAPS-24 HSDIP; ICAPS-24 Demo Track
– Developed, and submitted to IPC-2023, the first-ever PDDL representation for the Rubik's Cube domain, a complex combinatorial puzzle, showcasing the potential of symbolic representations in automated planning.

**Power State Identification through Non-Intrusive Harmonics Analysis** — Jan 2022–May 2022
*Related Publications:* IAAI/AAAI-23
– Formulated the State Identification Problem (SIP), a novel approach to discern unique power usage patterns through harmonics data in real-world scenarios.
– Conducted in-depth data analysis on electricity data, encompassing a comprehensive set of 221 features, spanning across key sectors: manufacturing, education, and hospital, to derive insights into intricate power consumption patterns.

## 📄 Selected Publications

- **Bharath Muppasani**, Protik Nag, Biplav Srivastava, Vignesh Narayanan. *On Generalized Planning for Controlling Opinion Networks: Interpreting Human-AI Dialog States and Beliefs.* Generalization in Planning (GenPlan) Workshop, AAAI (2025).
- **Bharath Muppasani**, Vignesh Narayanan, Biplav Srivastava and Michael N. Huhns. *Towards Effective Planning Strategies for Dynamic Opinion Networks.* NeurIPS Main Track (2024)
- **Bharath Muppasani**, Nitin Gupta, Vishal Pallagani, Biplav Srivastava, Raghava Mutharaju, Michael N. Huhns and Vignesh Narayanan. *Building a Plan Ontology to Represent and Exploit Planning Knowledge and Its Applications.* CODS-COMAD (2024).
- **Bharath Muppasani**, Vignesh Narayanan, Biplav Srivastava and Michael N. Huhns. *Expressive and Flexible Simulation of Information Spread Strategies in Social Networks Using Planning.* AAAI - Demonstrations Track (2024) 🏆 **Best Demo Award**.
- **Bharath Muppasani**, Vishal Pallagani, Biplav Srivastava, Forest Agostinelli. *Comparing Rubik's Cube Solvability in Domain-Independent Planners Using Standard Planning Representations for Insights and Synergy with Upcoming Learning Methods.* Heuristics and Search for Domain-Independent Planning (HSDIP) Workshop at ICAPS (2024).
- Vishal Pallagani, **Bharath Muppasani**, Kaushik Roy, Francesco Fabiano, Andrea Loreggia, Keerthiram Murugesan, Biplav Srivastava, Francesca Rossi, Lior Horesh, Amit Sheth. *On the Prospects of Incorporating Large Language Models (LLMs) in Automated Planning and Scheduling (APS).* ICAPS (2024).
- Vishal Pallagani, **Bharath Muppasani**, Biplav Srivastava, Francesca Rossi, Lior Horesh, Keerthiram Murugesan, Andrea Loreggia, Francesco Fabiano, Rony Joseph, Yathin Kethepalli. *Plansformer Tool: Demonstrating Generation of Symbolic Plans Using Transformers.* IJCAI Demonstrations Track (2023)
- Kausik Lakkaraju, Sai Krishna Revanth Vuruma, Vishal Pallagani, **Bharath Muppasani** and Biplav Srivastava. *Can LLMs be Good Financial Advisors?: An Initial Study in Personal Decision Making for Optimized Outcomes.* FinPlan Workshop at ICAPS (2023).
- **Bharath Muppasani**, Cheyyur Jaya Anand, Chinmayi Appajigowda, Biplav Srivastava and Lokesh Johri. *A Dataset and Baseline Approach for Identifying Usage States from Non-intrusive Power Sensing with MiDAS IoT-Based Sensors.* IAAI Technical Track on deployed Highly Innovative Applications of AI, AAAI (2023)
- **Bharath Muppasani**, Vishal Pallagani, Kausik Lakkaraju, Shuge Lei, Biplav Srivastava, Brett Robertson, Andrea Hickerson and Vignesh Narayanan. *On safe and usable chatbots for promoting voter participation.* AI Magazine, Volume 44, Issue 3, AAAI (2023)

## 💼 Experience

**Research Intern** — **Kharagpur, India**
*Indian Institute of Technology, Kharagpur* — May 2018–Aug 2018
*Advisor:* Niloy Ganguly, *Mentor:* Madhumita Mallick — *Smart Home Automation*

- Conducted in-depth analysis on multi-resident smart home data from Samsung, focusing on time, sensor, and frequency-based segmentation. Identified challenges such as the absence of activity annotations and the complexities of classifying activities.
- Initiated the groundwork for the development of an unsupervised algorithm aimed at labeling and classifying activities within smart homes. Proposed innovative measures to verify the goodness of clusters, setting the stage for future research and improvements in smart home data analysis and activity classification.
- Recognized by Samsung for outstanding research contributions in the field of smart home data analysis.

## 📄 Patents

- **Improving Planner Performance by Learning and Using Metadata of Experiences**
  *Accepted on March 24, 2023*
  *Patent Office: United States (US)*
  *Inventors:* Vishal Pallagani, **Bharath Muppasani**, Biplav Srivastava

in niting9

# NITIN GUPTA

+1 803 470 9214 | @ niting1209@gmail.com | g-nitin | Columbia, SC

Portfolio


## EDUCATION

**University of South Carolina** **GPA: 4.0/4.0**

***M.S.**: Computer Science; **Concentration**: Artificial Intelligence* *Aug 2025 – May 2026*

***B.S.**: Computer Science; **Minors**: Mathematics, Data Science* *Aug 2024 - May 2026*

***Relevant coursework:*** Neuromorphic & Edge Computing**,** Machine Learning Systems**,** Artificial Intelligence**,** Big Data Analytics**,** Software Engineering**,** Computer Architecture**,** Database Design

## SKILLS

**Advanced** in Python, PyTorch, Sklearn, CUDA, Java, C++, Linux, git, LaTeX — **Native** in English, Hindi, Punjabi

**Intermediate** in JavaScript, React, SQL, NoSQL, R, SCRUM, Apache Hadoop — **Proficient** in Spanish

## EXPERIENCE

**AI Researcher** — Columbia, SC

*AI Institute of USC, University of South Carolina* — *Aug 2024 – Present*

- Actively conducting research at the intersection of Neuro-symbolic AI and Natural Language Processing
- Developed innovative chatbot systems to promote voter participation among vulnerable groups
- Developed automated literature categorization tools, contributing to advancements in AI applications

**Volunteering Ambassador** — Columbia, SC

*Trew Friends* — *Aug 2021 – Present*

- Promoted organ, eye, and tissue donation, informing and educating people about the importance of donation.
- Successfully persuaded over 300 individuals to register as donors on the national organ donor registry.

**Explainable AI Research Intern** — Columbia, SC

*AI Institute of USC* — *May 2024 – Aug 2024*

- Developed a framework using Natural Language Processing and enriched knowledge graphs to generate transparent, context-aware explanations for AI-generated plans
- Conducted traffic data analysis for the SCDEC by identifying collision patterns and evaluating safety programs

**Machine Learning Researcher** — Columbia, SC

*Digital Research Services, University of South Carolina* — *Aug 2022 – Apr 2024*

- Conducted data mining on South Carolina Laws using Python and Natural Language Processing
- Trained machine-learning algorithms to study the prepared session laws and identify Jim Crow Language

**Material Analysis Researcher** — Columbia, SC

*Jefferson Lab, University of South Carolina* — *Jun 2022 – Oct 2022*

- Enhanced high energy physics research by developing C++ analysis code to process simulated data for various magnet materials, leading to potential cost savings of over $1,000,000

## PUBLICATIONS

**Towards Enhancing Road Safety in South Carolina Using Insights from Traffic and Driver-Education Data** — *AAAI 2025*
*N. Gupta, B. Muppasani, et al.*

**Revisiting LLMs in Planning from Literature Review: a Semi-Automated Analysis Approach and Evolving Categories Representing Shifting Perspectives** — *ICAPS 2025 (Accepted, Pending Archival)*
*V. Pallagani, N. Gupta, B. Muppasani, B. Srivastava*

**Building a Plan Ontology to Represent and Exploit Planning Knowledge and Its Applications** — *Discover Data Journal (Pending Review)*
*B. Muppasani, N. Gupta, V. Pallagani, B. Srivastava, et al.*

## AWARDS & ACHIEVEMENTS

**AAAI Student Scholar** encouraging student participation in the AI research community — *Feb 2025*

**McNAIR Junior Fellowship** for undergraduate computer science research. — *May 2024*

**Phi Beta Kappa Freshman Award** for outstanding academic performance at USC. — *Apr 2023*

## PERSONAL PROJECTS

**ArtEdge: Real-Time Neural Style Transfer for Mobile Devices** — ArtEdge
- An iOS application performing Neural Style Transfer (NST) directly on-device using edge computing

**Beacon Of Hope: A Personalized Meal Plan Recommender** — BOH
- A personalized, machine learning based, meal plan recommender system providing *tailored nutrition recommendations* & *visualizations*

**RoostAI: A University Centered Chatbot** — RoostAI
- A deployed conversational RAG-based LLM system, allowing USC student to *easily explore campus services*

# CV of Biplav Srivastava

## Employment History

- 2020-: Professor, AI Institute, University of South Carolina, Columbia, SC, USA
- 2001-2010: IBM
  - 2018-2020: Distinguished Data Scientist, Master Inventor IBM Chief Analytics Office, NY, USA
  - 2016-2018: Research Staff Member, Master Inventor, IBM T.J. Watson Lab, NY, USA
  - 2009-2016: Senior Researcher, Master Inventor IBM Research, New Delhi, India
  - 2007-2009: Research Staff Member, IBM T.J. Watson Lab, NY, USA
  - 2001-2007: Research Staff Member IBM Research, Delhi & Bangalore, India
- 2000-2001: Core Technology Architect, R&D Bodha.com, Los Altos, CA, USA
- 1996-2000: Senior Staff Engineer, Philips Semi. (VLSI Tech.), Phoenix, USA; PhD Scholar, Arizona State University, USA

## Education

- Ph.D.: Arizona State University, Tempe, AZ, USA, 2000. Dissertation Topic: Efficient Planning by Effective Resource Reasoning
- M.S.: Arizona State University, Tempe, AZ, USA, 1996. GPA:4.0/4.0
- B.Tech.: Indian Institute of Technology (IIT), B.H.U., India, 1993. GPA:9.06/10
- Micro-MBA: Management module at IBM taught by faculties of top-US MBA schools.

## Research Interests

He is interested in enabling people to make rational decisions despite real world complexities of poor data, changing goals and limited resources by augmenting their cognitive limitations with technology. His expertise is in Artificial Intelligence, Sustainability and Services, with proven track record of many science firsts and high-impact innovation ($B+) in global business environment. His research initiatives are on promoting goal-oriented, ethical, human-machine collaboration via natural interfaces using domain and user models, learning and planning. In particular, he has been exploring the role of reasoning (planning) and learning for human-agent dialogs in individual and group settings.

Biplav has worked extensively in planning starting with initial research in unifying classical planning methods (**?**) and synthesizing planners from specifications (**??**). His PhD was a fundamental work in teasing causal and resource reasoning in planning with RealPlan (**???**). At IBM Research, he initiated work on diverse planning with collaborators in academia, which was the first in the field on how to measure plans based on a variety of distance metrics and automatically generate a set of plans with quality constraints (e.g., dis-similarity) in the presence of incomplete or unknown preferences. He formulated the problem based on experience of applying planning in automated system management (i.e., web services composition and autonomic computing), and it lead to work published in IJCAI (**??**) and AIJ (**?**), and a PhD thesis. The work has wide applications to the original motivations, as well as areas like traffic management and security.

In applications of planning to web services, Biplav lead the influential end-to-end automated web service composition approach, called Synthy, that composes SOA-compliant components (Web Services) to create a new composite web service while meeting the desired functional and non-functional requirements (**??????**). Synthy is as much significant for tackling various aspects of web service composition as for applying different facets of computer science – AI Planning, Web Semantics, Distributed Computing and Optimization. Another first in web services is Semaplan, an approach which Biplav co-invented to represent and match services using semantics and planning techniques (**?**) that was implemented and made available as a major feature in IBM's software offering. In autonomic computing, Biplav articulated the role of declarative methods like planning in automated control, incorporated a Java-based family of planners in agent toolkit (ABLE) and demonstrated their real-world usage with selective automation ((**??**)).

Biplav's current focus in planning is neuro-symbolic methods and generalized planning.

## Recognition and Awards

- Faculty Awards from Cisco Research (2022) and JP Morgan Chase (2023)
- AAAI Senior Member, recognized by Association for Advancement of Artificial Intelligence for significant contributions and service to AI community (Feb 2020)
- AAAS Leshner Public Engagement Fellow, 2020-2021 cohort of the American Association for the Advancement of Science Leshner Leadership Institute for Public Engagement with Science, focusing on AI (Feb 2020)
- IBM Master Inventor – for outstanding patent activities (2011-2017, 2019-2022)
- Represented IBM at Partnership on AI's Working Group on Fair, Transparent and Accountable AI (2018-2020). https://www.partnershiponai.org
- IEEE Senior Member, for technical achievements (Dec 2018)

**IDENTIFYING INFORMATION**:

NAME: Narayanan, Vignesh

POSITION TITLE: Assistant Professor, AI Institute

PRIMARY ORGANIZATION AND LOCATION: University of South Carolina, Columbia, SC, United States

**Professional Preparation**:

| ORGANIZATION AND LOCATION | DEGREE (if applicable) | RECEIPT DATE | FIELD OF STUDY |
|---|---|---|---|
| Washington University in St. Louis, St. Louis, Missouri, United States | Postdoctoral Fellow | 09/2017 - 07/2021 | Applied mathematics and Computational neuroscience |
| Missouri University of Science and Technology, Rolla, Missouri, United States | PHD | 07/2017 | Electrical engineering, Control systems |
| National Institute of Technology, Kurukshetra, Haryana, India | MOTH | 07/2014 | Control systems |
| SASTRA University, Thanjavur, Tamilnadu, India | BOTH | 07/2012 | Electrical and Electronics Engineering |

**Appointments and Positions**

2021 - present Assistant Professor, AI Institute, University of South Carolina, Columbia, SC, United States

2021 - present Assistant Professor, Department of Computer Science and Engineering, University of South Carolina, Columbia, SC, USA

2021 - present Core Associate, Carolina Autism and Neurodevelopment Research Center, University of South Carolina, Columbia, SC, United States

2017 - 2021 Postdoctoral Fellow, Washington University in St. Louis, St. Louis, MO, USA

**Products**

*Products Most Closely Related to the Proposed Project*

1. Narayanan V, Zhang W, Li J. Duality of Ensemble Systems through Moment Representations. IEEE Transactions on Automatic Control. 2024 October; 69(10). Available from: 10.1109/TAC.2024.3397159
2. Narayanan V, Bordoh LK, Kiss IZ, Li JS. Inferring networks of chemical reactions by curvature analysis of kinetic trajectories. Phys Chem Chem Phys. 2025 May 14;27(19):9962-9969. PubMed Central PMCID: PMC12160075.
3. Miao W, Narayanan V, Li JS. Parallel residual projection: a new paradigm for solving linear inverse problems. Sci Rep. 2020 Jul 30;10(1):12846. PubMed Central PMCID: PMC7393146.
4. Miao W, Narayanan V, Li JS. Interpretable Design of Reservoir Computing Networks Using Realization Theory. IEEE Trans Neural Netw Learn Syst. 2023 Sep;34(9):6379-6389. PubMed

Central PMCID: PMC10567084.

5. Narayanan V, Li JS, Ching S. Biophysically interpretable inference of single neuron dynamics. J Comput Neurosci. 2019 Aug;47(1):61-76. PubMed PMID: 31468241.

*Other Significant Products, Whether or Not Related to the Proposed Project*

1. Zhang W, Narayanan V, Li S. Robust Population Transfer for Coupled Spin Ensembles. 2019 IEEE 58th Conference on Decision and Control (CDC). 2019; :419-424. isbn: 1728113989
2. Yu YC, Narayanan V, Li JS. Moment-Based Reinforcement Learning for Ensemble Control. IEEE Trans Neural Netw Learn Syst. 2024 Sep;35(9):12653-12664. PubMed Central PMCID: PMC10676148.
3. Zhang W, Narayanan V, Li J. Dynamic Function Learning through Control of Ensemble Systems. When Machine Learning meets Dynamical Systems: Theory and Applications (MLmDS), AAAI 2023. 2023.
4. Narayanan V, Modares H, Jagannathan S, Lewis FL. Event-Driven Off-Policy Reinforcement Learning for Control of Interconnected Systems. IEEE Trans Cybern. 2022 Mar;52(3):1936-1946. PubMed PMID: 32639933.
5. Muppasani B, Nag P, Narayanan V, Srivastava B, Huhns M. Towards Effective Planning Strategies for Dynamic Opinion Networks. Advances in Neural Information Processing Systems. 2024; 37:137046-137104.

**Certification:**

I certify that the information provided is current, accurate, and complete. This includes but is not limited to information related to domestic and foreign appointments and positions.

I also certify that, at the time of submission, I am not a party to a malign foreign talent recruitment program.

Misrepresentations and/or omissions may be subject to prosecution and liability pursuant to, but not limited to, 18 U.S.C. §§ 287, 1001, 1031 and 31 U.S.C. §§ 3729-3733 and 3802.

Certified by Narayanan, Vignesh in SciENcv on 2025-07-16 10:35:01