Artificial Intellligence: An Introduction

Wei Wang(王伟)

Engineering Research Centre of Applied Technology on Machine Translation and Artificial Intelligence , Macao Polytechnic University

weiwang@mpu.edu.mo; 匯智樓 (WUI CHI)-5/F, N56

Jan. 12, 2026

Outline

- **AI Models**

AI Models

AI models refer to different approaches or frameworks that are used to represent and solve problems in the field of AI.

These models provide a structured way to understand and analyze complex systems and make intelligent decisions.

Reflex

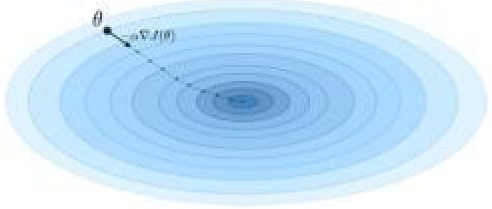

- Feature vector, score, margin, residual
- Linear predictors, non-linear predictors
- Loss minimization, stochastic gradient descent, backpropagation, regularization

PDF Web

States

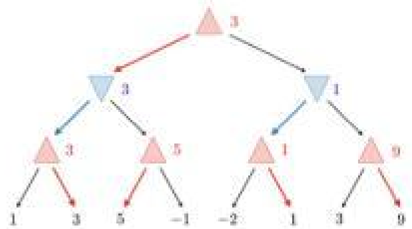

- Tree search, graph search, A star search
- Learning costs, relaxation
- Markov decision processes, game playing

PDF Web

Variables

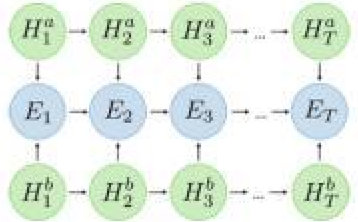

- Constraint satisfaction problem, backtracking search
- Beam search, Gibbs sampling
- Bayesian networks, probabilistic program, inference

PDF Web

Logic

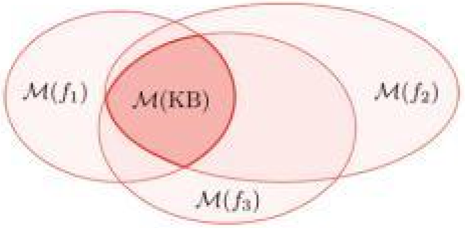

- Notations, knowledge base, forward inference, rule properties
- Propositional logic, modus ponens, resolution rule, conjunctive normal form
- First-order logic, substitution, unificaion

PDF Web

Source: https://stanford.edu/~shervine/teaching/cs–221/

AI models define how an agent perceives, reasons, and acts.

Different models suit different environments and tasks.

AI Models

1. Logic-based models

- Symbolic representation of classes of objects.
- Deductive Reasoning.
- **Apps**: Question Answering Systems, Natural Language Understanding, Expert system
- **Options**: Propositional Logic , First-Order Logic, Knowledge Base.

3. Variables-based models (Uncertainty)

- Solution in an assignment of values for a set of variables.
- **Apps**: Soduko, Speech Recognition, and Face Recognition.
- **Options**: Convolutional Neural Networks, Bayesian Networks, Constraint Satisfaction, Factor Graphs, and Dynamic Ordering.

2. States-based models

- Solutions are defined as a sequence of steps.
- Model a task as a graph of states and a solution as a path in the graph.
- A state captures all of the relevant information about the past in order to act in the future.
- **Apps**: Navigation and Games.
- **Options**: Tree Search (Breadth-first, Depth-first search), Graph search (Dynamic programming), Markov decision processes, Hidden Markov Models, Game playing

4. Reflex-based models

- Given a set of <Input, Output> pairs of training data, learn a set of parameters that will map input to output for future data.
- **Apps**: Classification and Regression.
- **Options**: Artificial Neural Networks (ANN), Decision Trees, Support Vector Machines, Regression, Principal Component Analysis, K-Means Clustering, and K-Nearest Neighbor

AI Models

Reflex-based Models

A reflex-based model simply performs a fixed sequence of computations on a given input. Examples include most models found in machine learning, from simple linear classifiers to artificial neural networks (ANN).

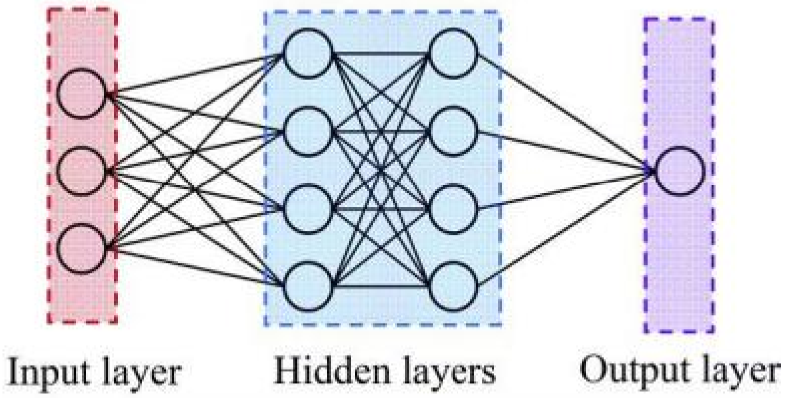

Direct mapping from percept to action;
No memory or internal state;
Fast but limited intelligence.

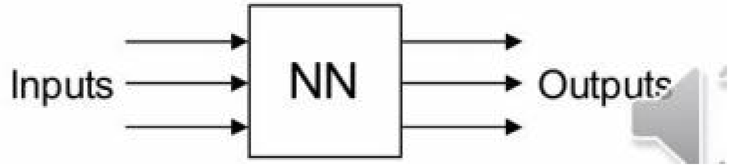

State-based Models

Reflex-based models are too simple for tasks that require more forethought (e.g., in playing chess or planning a big trip). State-based models overcome this limitation.

AI agents solve problems by moving between states using actions.

Many AI tasks can be framed as:

The world has a current situation (state)

The agent can take actions

Actions change the world into a new state

Goal: find a sequence of actions to reach a goal state

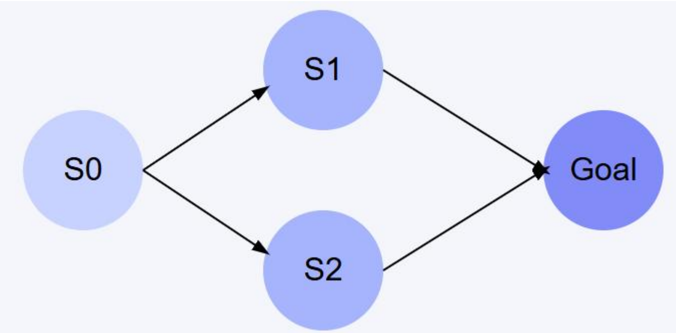

Used in:

Pathfinding (maps, robot navigation)

Puzzle solving (8-puzzle, Sudoku)

Automated planning

Game AI

1	2	3
5	6	0
7	8	4

1	2	3
5	8	6
0	7	4

AI Models

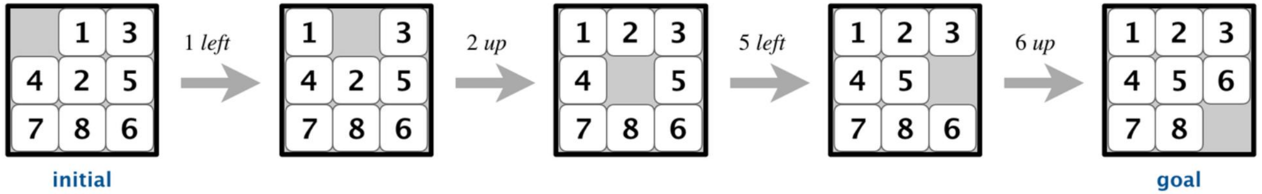

Core Components of a State-based Problem

A problem is defined by:

1.**State space**: all possible states

2.**Initial state**: where we start

3.**Actions / Operators**: what can be done

4.**Transition model**: how actions change states

5.**Goal test**: how to know we're done

6.**Path cost** (optional): measure of solution quality

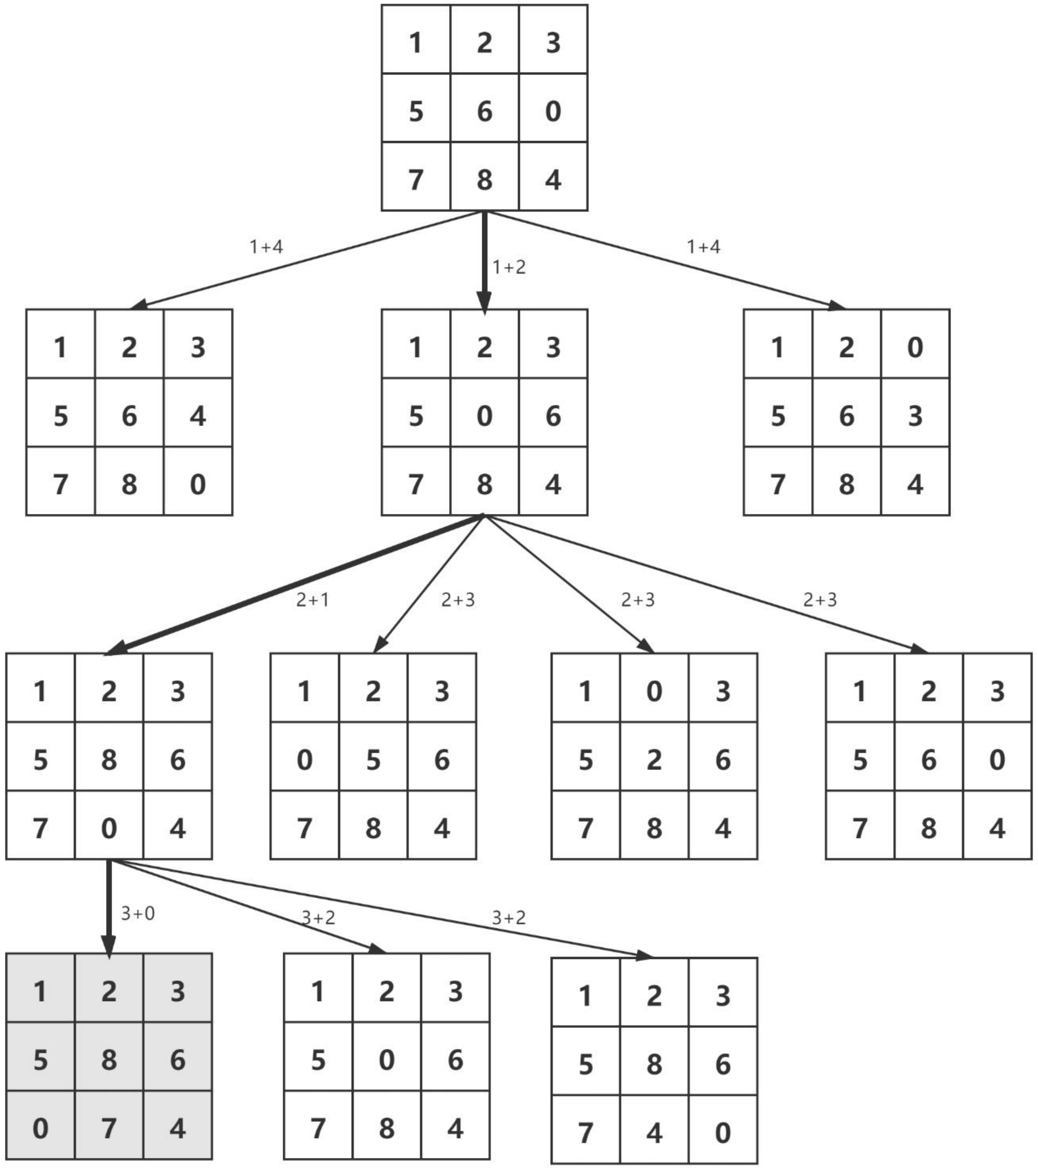

Uninformed Search (Blind Search)

Search without knowing where the goal is.
Common methods:

BFS (Breadth-First Search)： expand shallowest nodes first

DFS (Depth-First Search)： go deep first, backtrack later

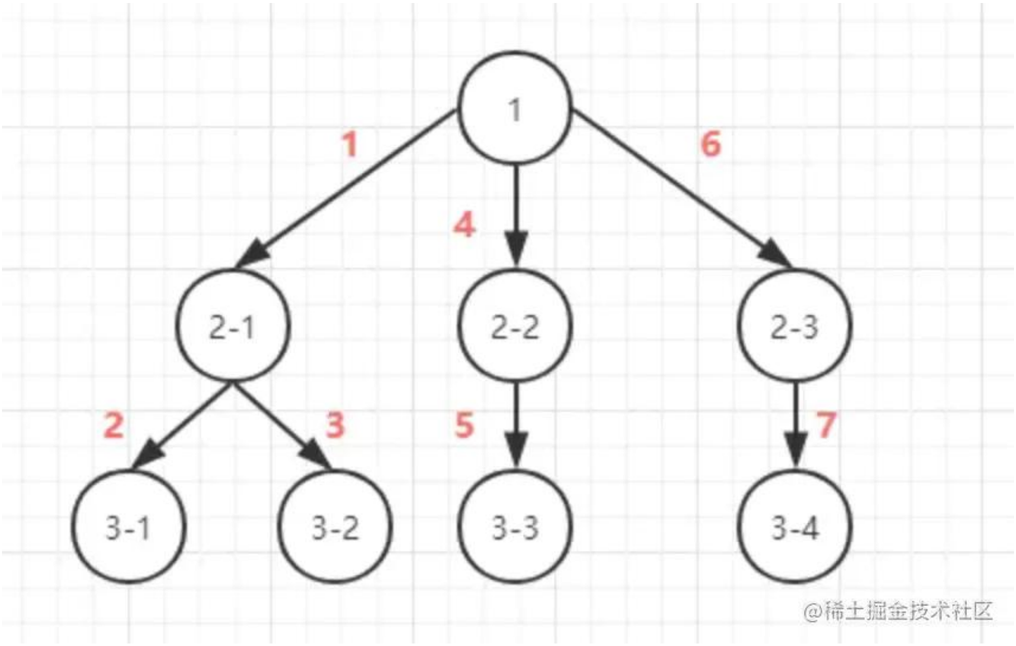

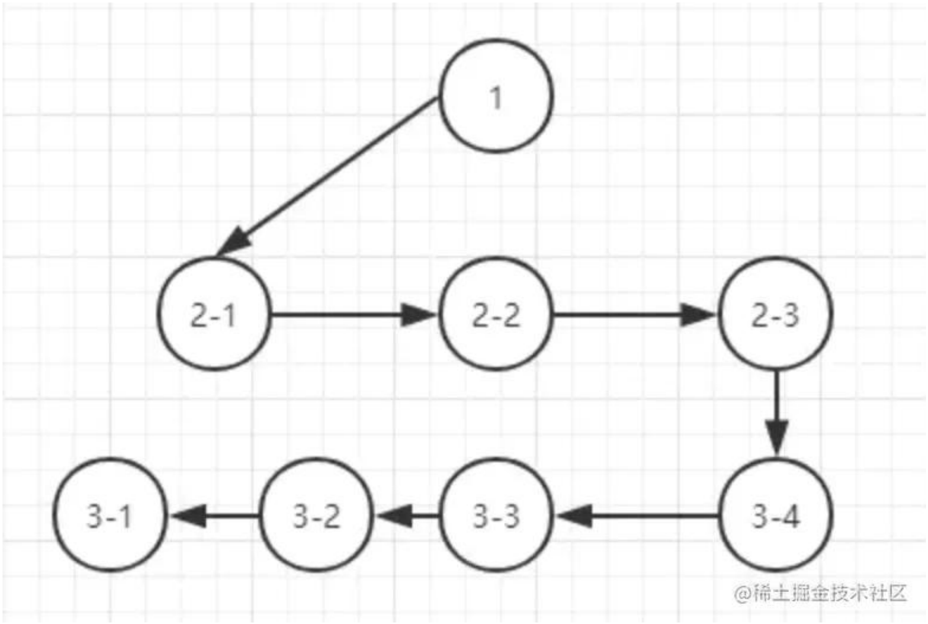

AI Models

Variables-based models

Logic-based models are a way to represent AI problems using formal logic (such as propositional logic or first-order logic), where:

Knowledge is written as **logical statements** (facts and rules)

Reasoning is done by applying inference rules

The system can derive **new conclusions** that logically follow from what it already knows

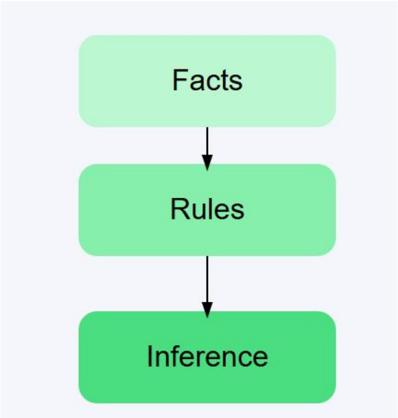

Logic-based models help AI:

Represent human knowledge clearly

Perform reasoning step-by-step

Explain why a conclusion is true

Support rule-based decision systems

Used in:

Expert systems (medical diagnosis, troubleshooting)

Knowledge graphs + reasoning

Legal reasoning

Automated theorem proving

Rule-based agents

AI Models

Components of a Logic-based Model

A typical logic reasoning system includes:

Symbols (objects, properties, relations)

Facts (what is known)

Rules (if–then relationships)

Queries (questions asked to the system)

Inference (reasoning process)

Advantages:

Clear and interpretable reasoning

Explainable decisions (“why true”)

Strong for structured knowledge and rules

Proven mathematical foundation

Challenges:

Hard to model uncertainty (probabilistic events)

Real-world knowledge can be incomplete/noisy

Large-scale reasoning may be expensive

Requires correct formal rules

AI Model Progression

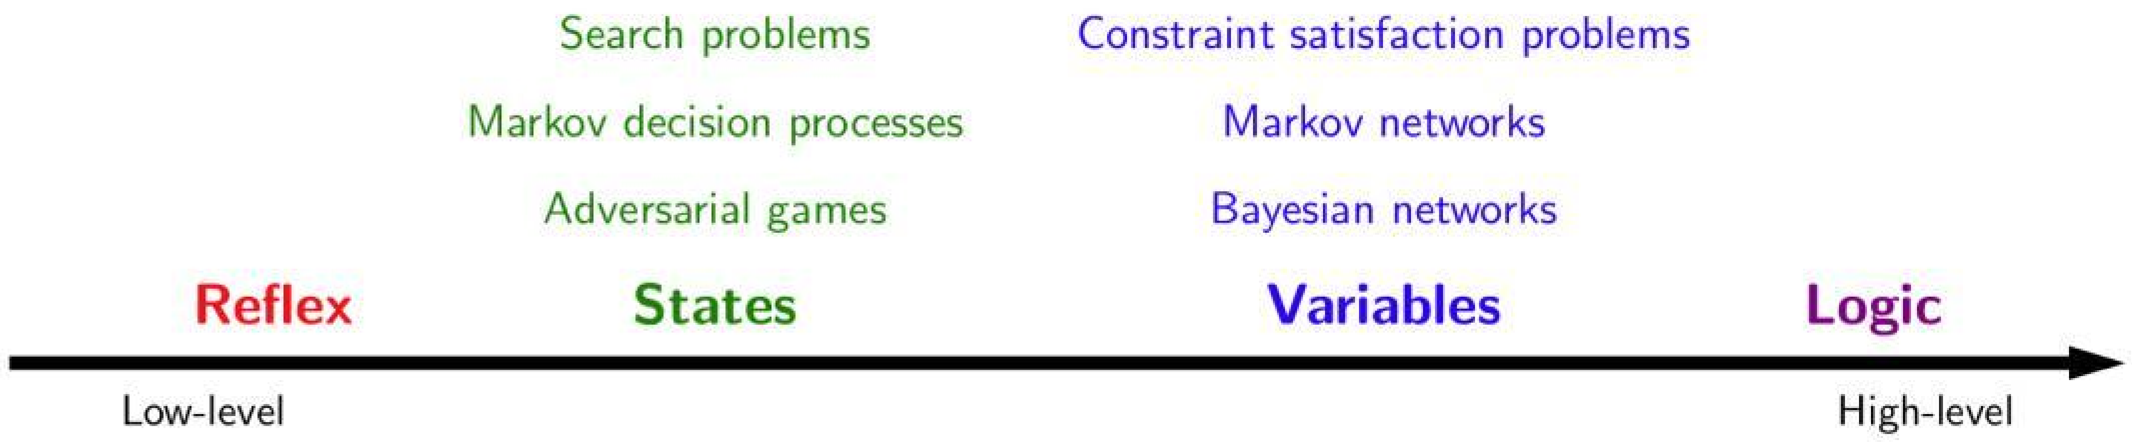

AI Models

The progression from low–level to high–level AI models reflects an evolution towards more complex and advanced capabilities. This progression enables increasing complexity, sophistication, memory, and reasoning capabilities in AI systems.

The incorporation of states and variables enables more sophisticated decision– making and capturing complex relationships. Additionally, logic–based AI models introduce logical representations and reasoning mechanisms, allowing for more advanced reasoning and handling of uncertainty.

AI Models

AI Model Progression

Reflex Models (Low-level AI): Simple and reactive, reflex models map inputs to outputs based on predefined rules or conditions. They lack internal states, making immediate responses suitable but lacking complex reasoning or memory.

States (Intermediate-level AI): Crucial in higher-level AI, states capture past inputs and actions, enabling historical context for sophisticated decision-making. They provide memory and track system evolution over time.

Variables (Intermediate-level AI): Important in both low-level and high-level AI, variables capture information and features in the problem domain. They serve as inputs, facilitating predictions and decisions. High-level AI models learn complex relationships between variables, enabling advanced reasoning.

Logic-based AI Models (High-level AI): These models emphasize reasoning and inference using logical representations. They encode knowledge through logical rules, deriving new information and handling uncertainty. Operating at a higher level of abstraction, they enable complex reasoning and rule-based decision- making.

AI Models

Questions

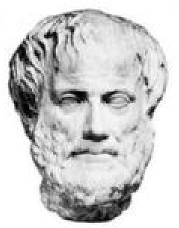

symbolic AI

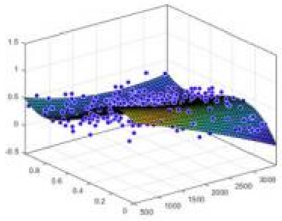

statistical AI

neural AI

There are three intellectual traditions in AI: Symbolic AI, Neural AI, and Statistical AI. Additionally, there are four AI models: Logic-based models, Variable-based models, State- based models, and Reflex-based models.

How do these AI models relate to the three intellectual traditions, and how can we assign the AI models to their respective traditions?

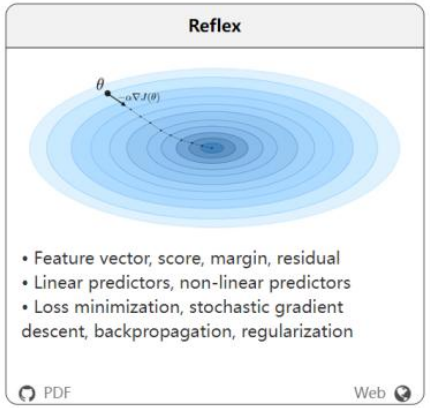

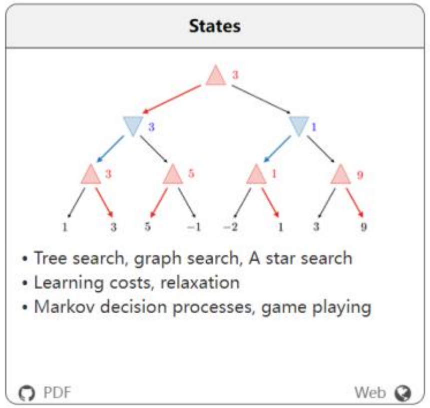

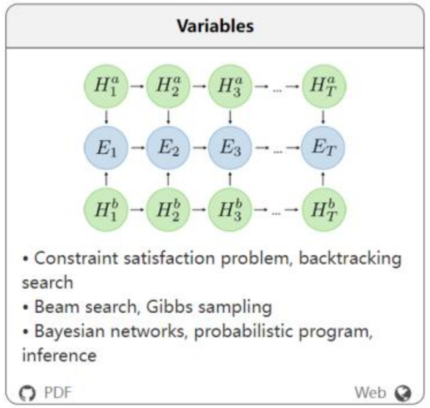

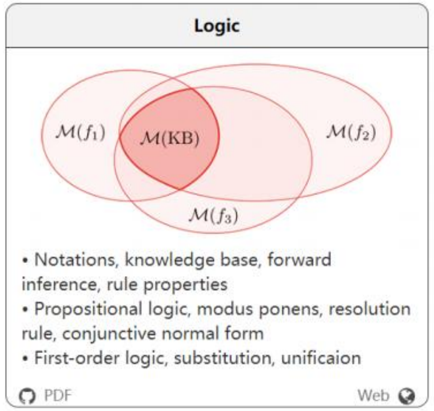

AI Models

The Answer

Symbolic AI: Symbolic AI is associated with logic-based models. Logic-based models, such as rule-based systems and expert systems, are rooted in symbolic AI. They use logical rules and formal reasoning to represent and manipulate knowledge.

Statistical AI: Statistical AI encompasses various statistical models, including reflex-based models and state-based models. Reflex-based models, such as linear regression and logistic regression, are used for data analysis and modeling by capturing relationships between variables. State-based models, such as Hidden Markov Models (HMMs) and Markov Decision Processes (MDPs), are employed for modeling sequential data and decision-making under uncertainty.

Neural AI: Neural AI is associated with variable-based models. Variable-based models, including neural networks, are a key component of neural AI. These models use mathematical algorithms and optimize the values of variables to learn and make predictions.

Outline

- AI Applications

AI Applications

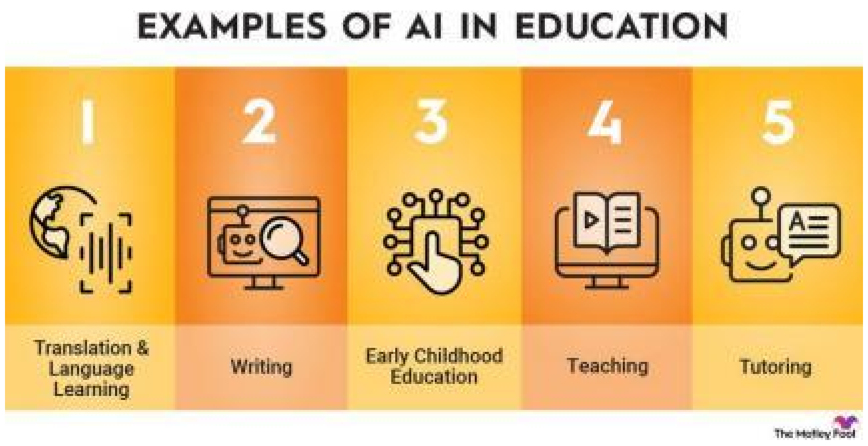

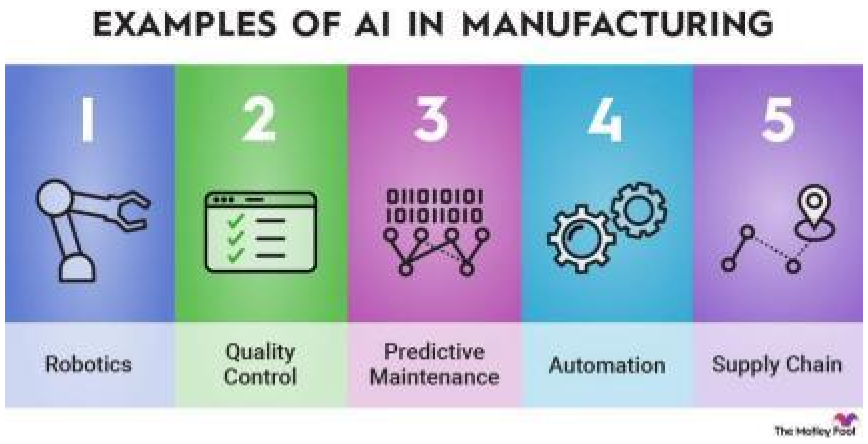

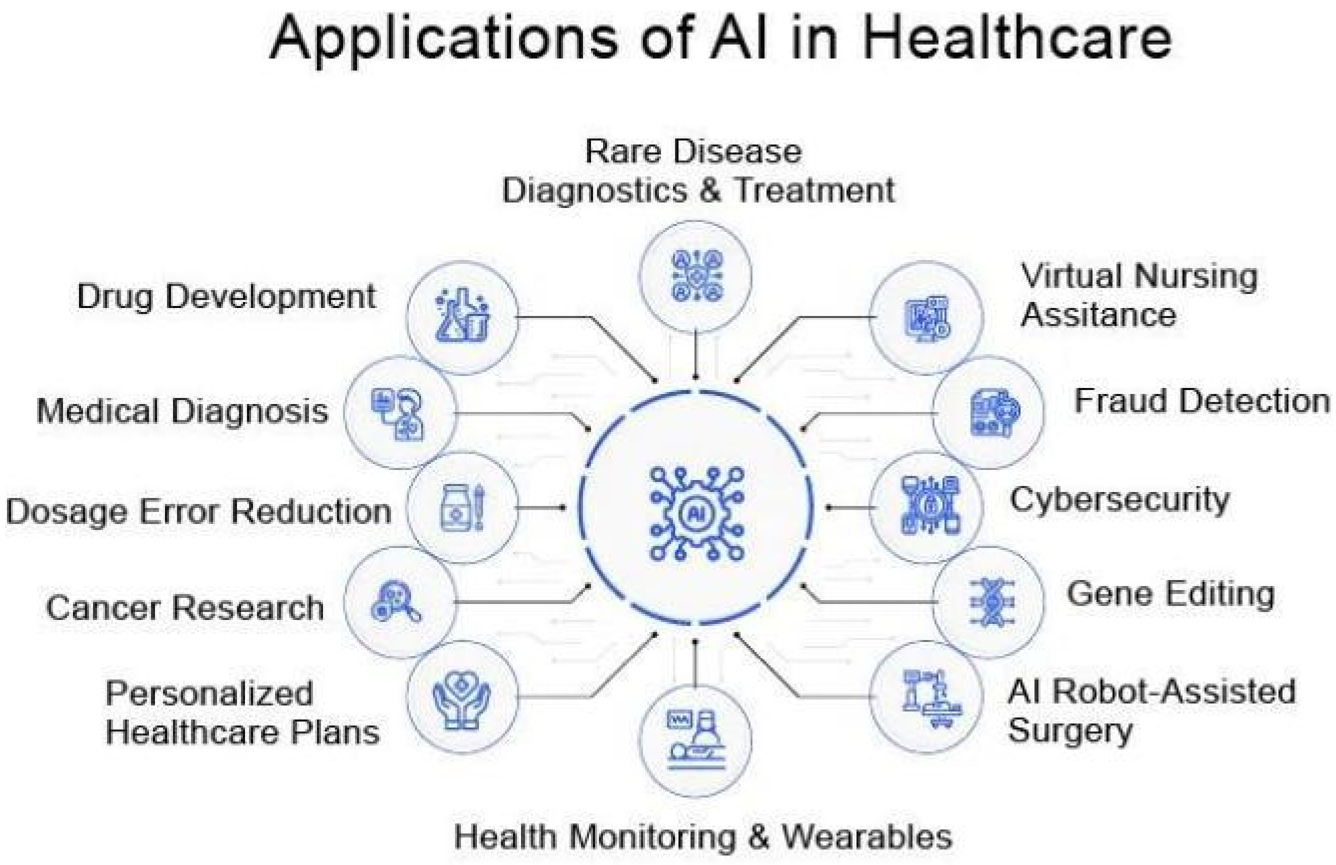

AI Applications

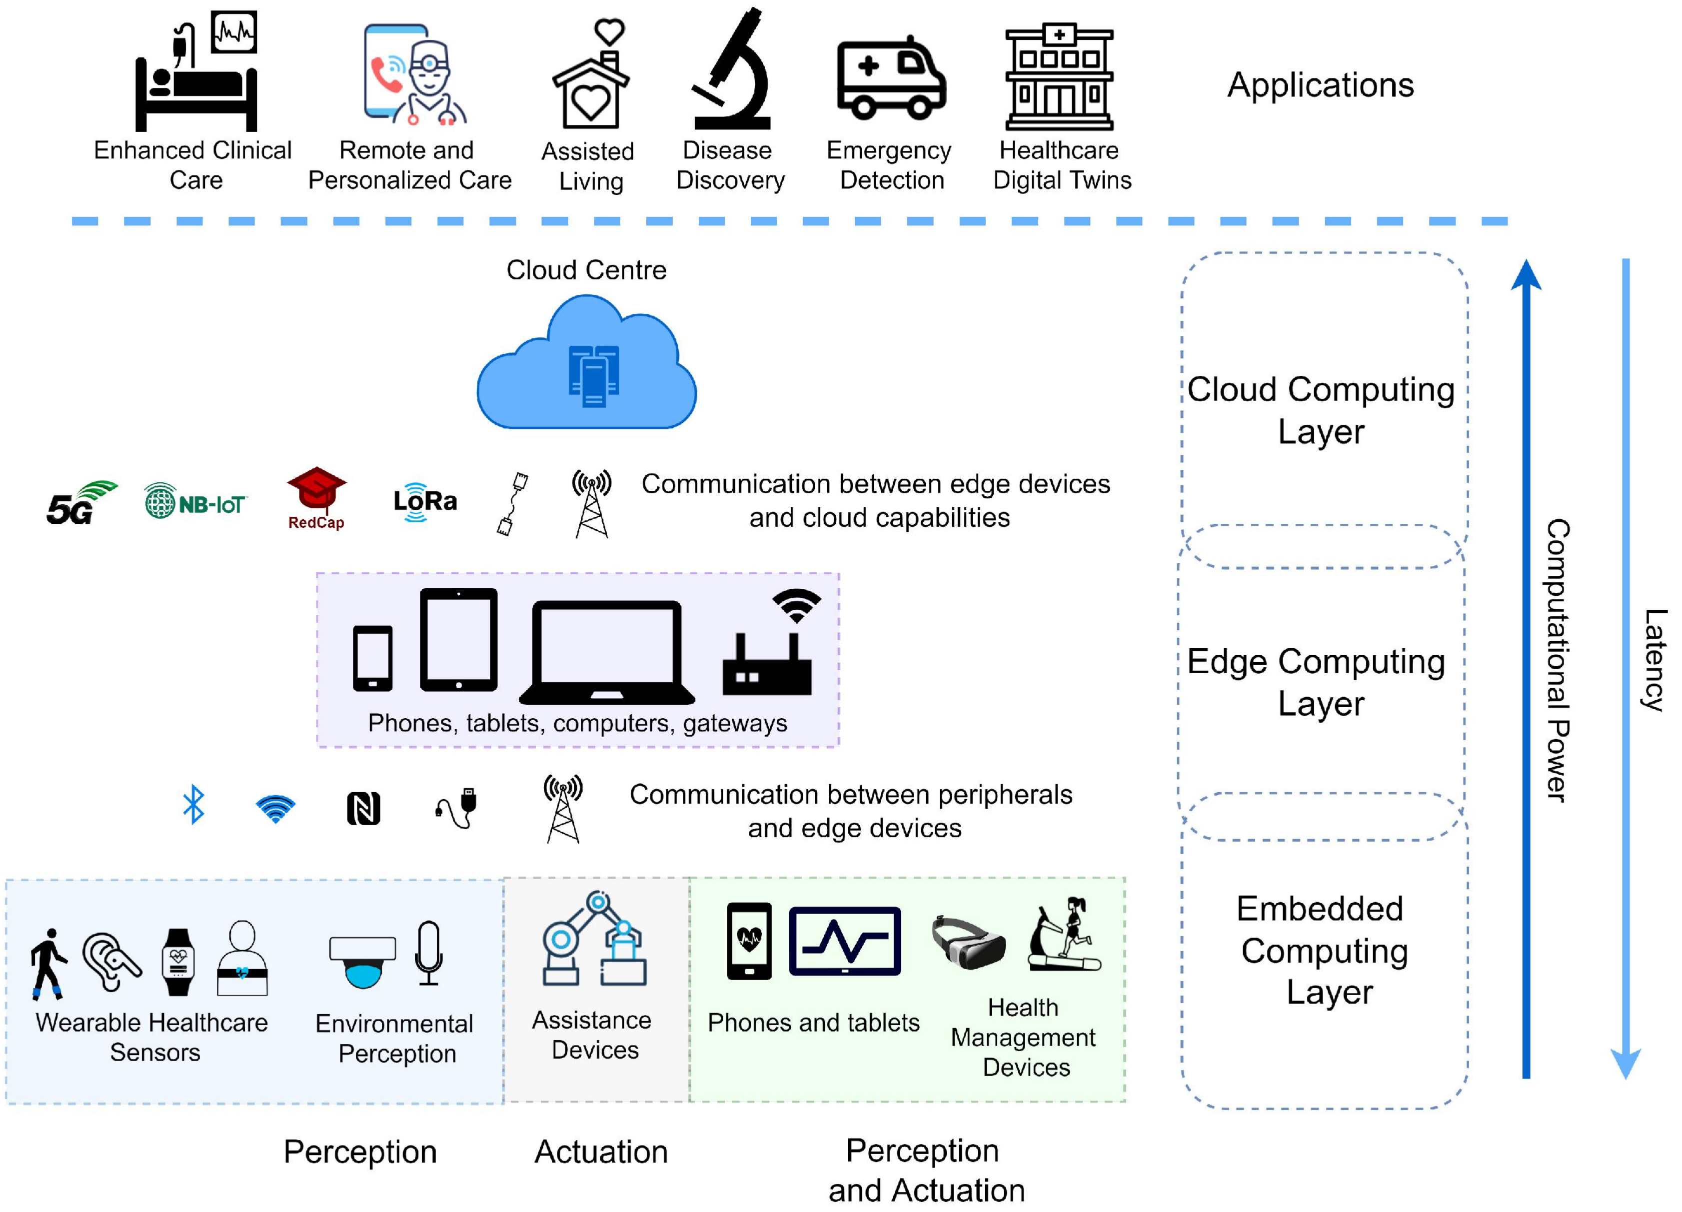

Mental Health

- Emotion recognition
- Depression recognition

Physical Health

- Heart Disease
- Alzheimer's disease

Baker, Stephanie, and Wei Xiang. "Artificial intelligence of things for smarter healthcare: a survey of advancements, challenges, and opportunities." IEEE Communications Surveys & Tutorials 25, no. 2 (2023): 1261-1293.

AI Applications

In Hospital: Auxiliary Diagnosis

- **Goal:** Precise and Optimized
- **Challenge:** High auality but Few and difficult to access

Out Hospital: Health Monitering

- **Goal:** Active amd Early
- **Challenge:** Expensive equipment, low accuracy

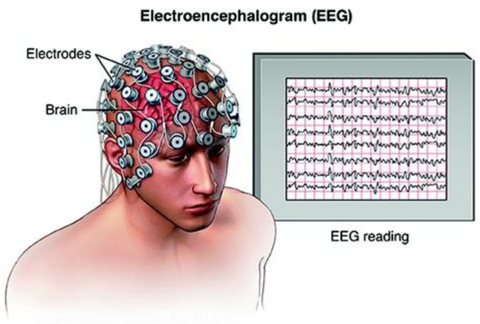

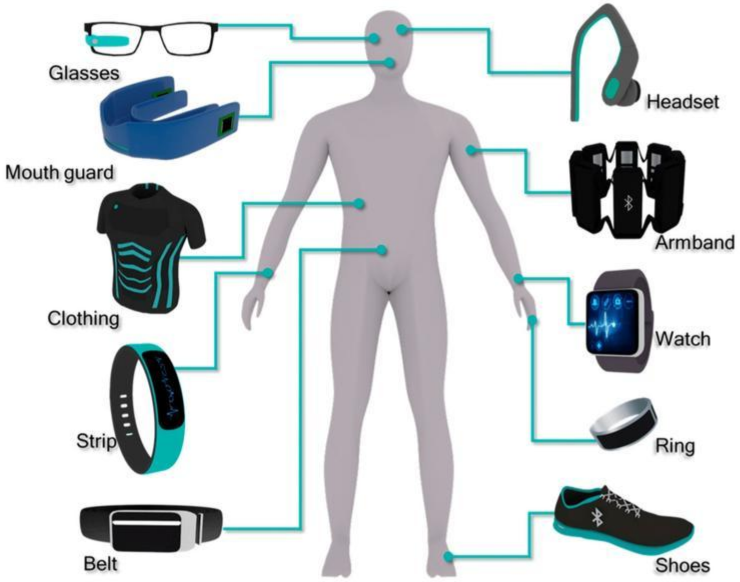

Fair, Accessible, and Systematic Health Services

--Healthy China 2030 Planning Outline

AI Applications

Single-lead ECG

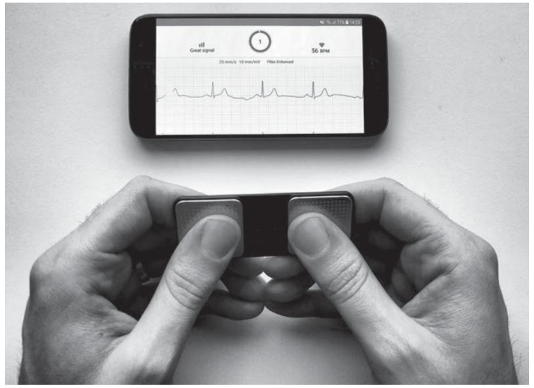

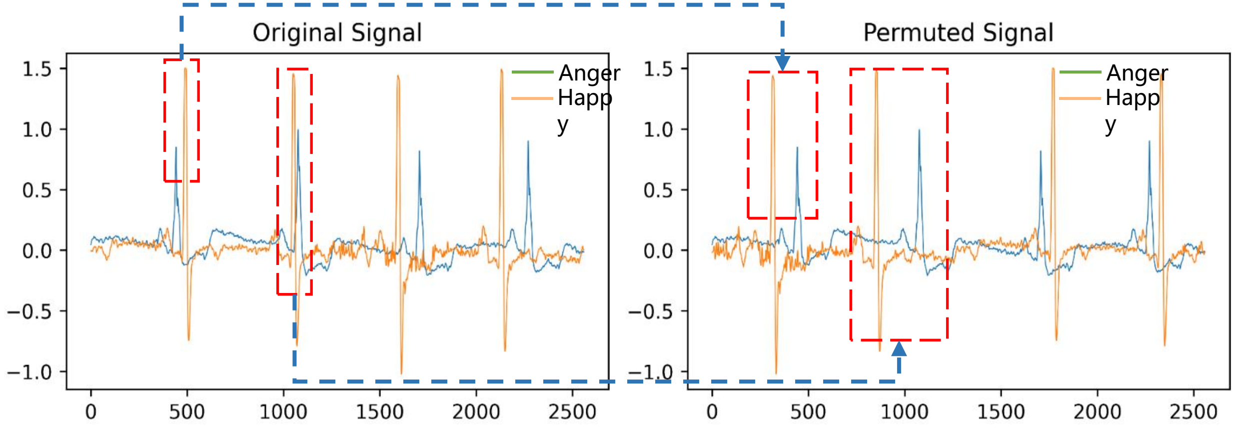

Challenge ❗

The signals of different tags are different and susceptible to individuality

Question ❓

How to effectively identify the difference of ECG signals under different categories

AI Applications

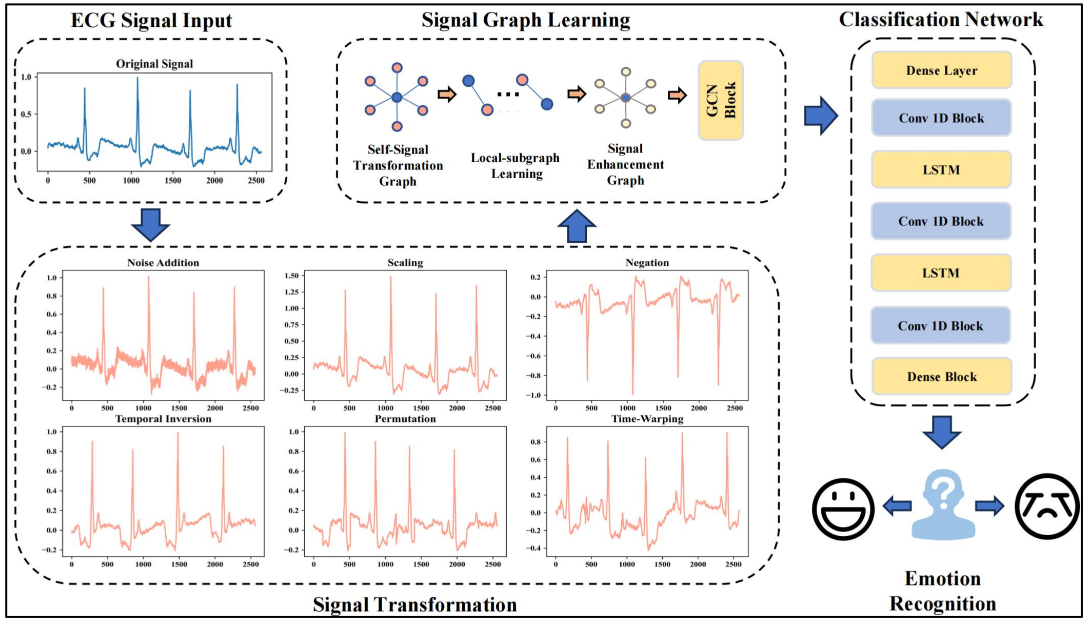

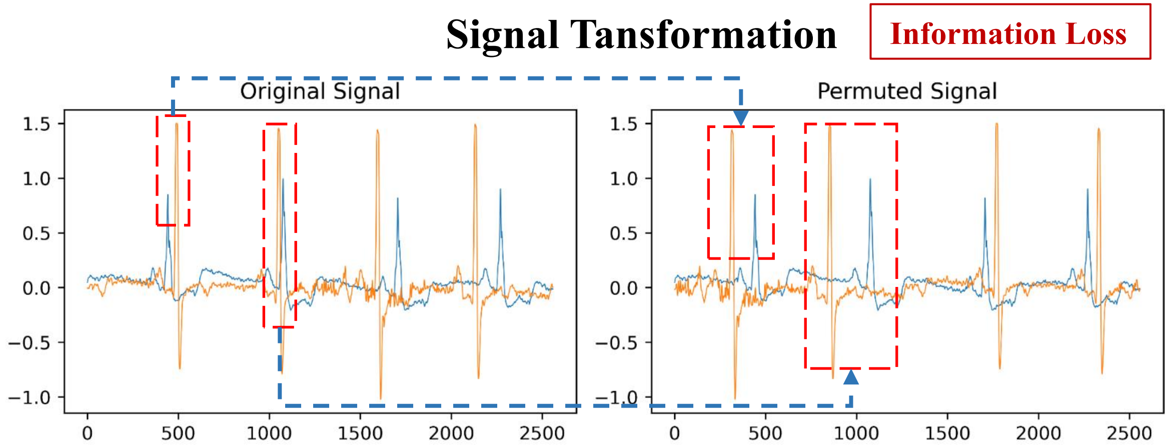

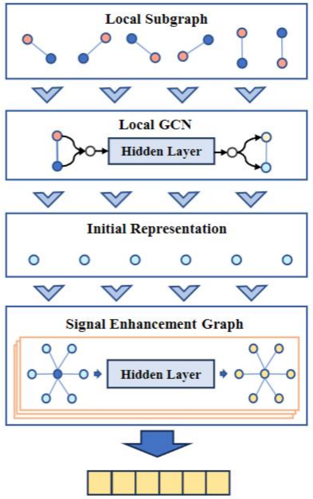

Graph Enhanced Representation

– **Target:** ***Recognize*** emotions via **single lead ECG data**

– **Innovations**:

- ✓ **Signal transformation** for signal enhancement;
- ✓ **Representation learning** for effective feature extraction and reducing information loss;

– **Application:** Intelligent interactive system based on convenient emotion recognition, i.e., mental health monitoring, user emotion monitoring, etc.

Experiments on WESAD Dataset

	DT	RF	Adaboost	LDA	KNN	MNSF	FSCNN	SSCNN	Inceptiontime	DeepCNN	GERECG
Precision	0.3886	0.6198	0.3889	0.2571	0.3563	0.6690	0.829	0.8565	0.9082	0.8380	**0.9111**
Recall	0.3846	0.4825	0.3853	0.2564	0.3072	0.6231	0.7421	0.8496	0.8945	0.7950	**0.9097**
F1	0.3857	0.4718	0.3862	0.2556	0.3058	0.6279	0.7279	0.8529	0.9068	0.8040	**0.9096**
Accuracy	0.4216	0.5844	0.4207	0.3036	0.3806	0.6824	0.8399	0.8702	0.9234	0.8563	**0.9240**

Removed	Precision	Recall	F1	Accuracy
Noise addition	0.8863	0.8574	0.8799	0.9050
Scaling	0.8866	0.8493	0.8684	0.9009
Negation	0.8855	0.8448	0.8601	0.9005
Temporal Inversion	0.8866	0.8566	0.8772	0.9058
Permutation	0.8920	0.8630	0.8834	0.9083
Time-warping	0.8853	0.8586	0.8760	0.8996

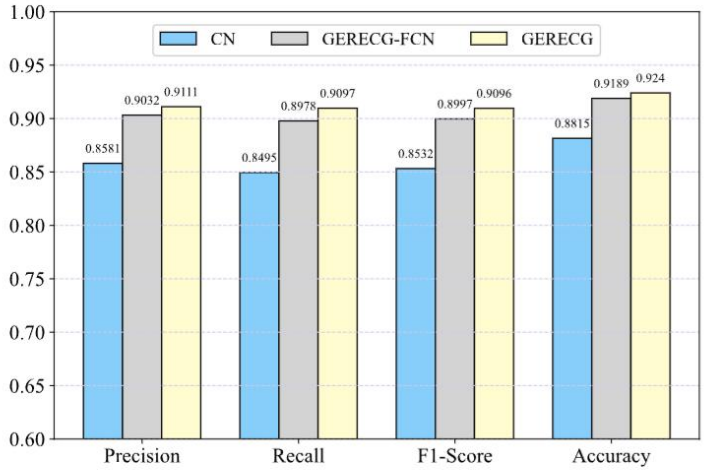

Chen, Jian, Yuzhu Hu, Lalit Garg, Thippa Reddy Gadekallu, Gautam Srivastava, and **Wei Wang***. "Graph enhanced low-resource ECG representation learning for emotion recognition based on wearable internet of things." **IEEE Internet of Things Journal,** vol. 11, no. 24, pp. 39056-39068 (2024). (**IF 8.9, JCR Q1**)

Modeling-Inference-Learning Paradigm

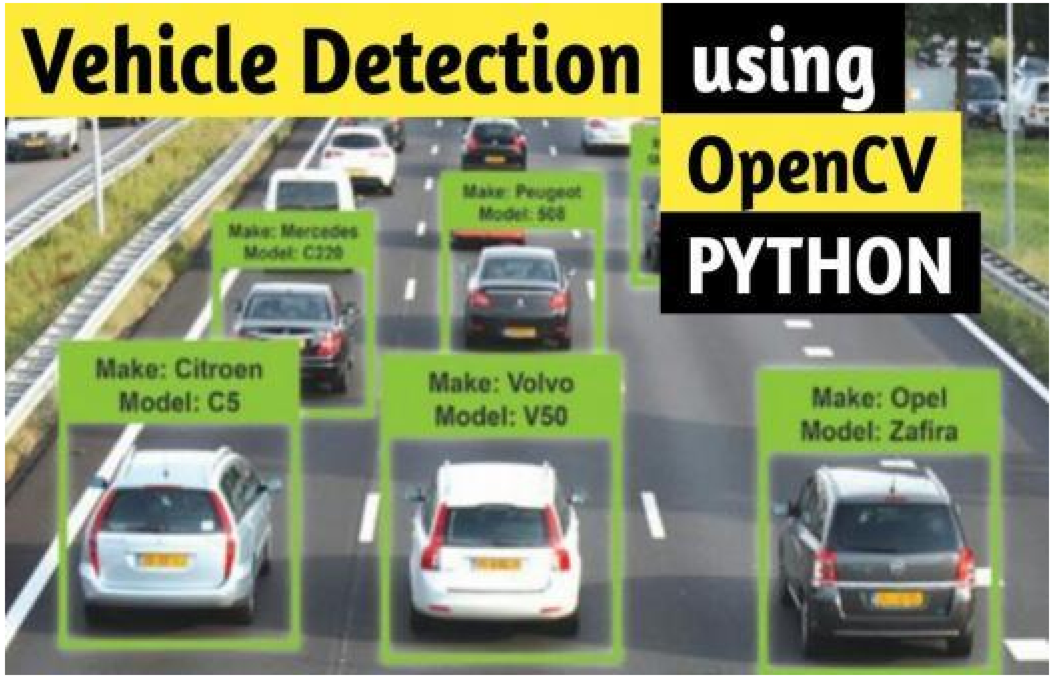

Modeling-Inference-Learning Paradigm

Modeling: This paradigm involves creating models that represent the problem domain or the system being studied. Models serve as abstractions of reality, capturing essential features and relationships. They can be created using various techniques such as rule-based systems, graphical models, or neural networks.

Inference: Inference refers to the process of reasoning and drawing conclusions based on available information or evidence. In this paradigm, AI systems use the models and available data to make logical deductions, probabilistic reasoning, or logical reasoning. Inference enables AI systems to make decisions, predict outcomes, or generate new knowledge.

Learning: Learning is a key paradigm in AI that focuses on systems acquiring knowledge or improving performance through experience or data. Machine learning algorithms, such as supervised learning, unsupervised learning, or reinforcement learning, are used to train AI systems and enable them to adapt, generalize, and make predictions based on new or unseen data.

Paradigm: Modeling

- The first pillar is modeling. Modeling takes messy real world problems and packages them into neat formal mathematical objects called **models**, which can be subject to rigorous analysis and can be operated on by computers.
- However, modeling is **lossy**: not all of the richness of the real world can be captured, and therefore there is an art of modeling: **what does one keep versus ignore**?
- (An exception to this are games such as Chess, Go, or Sudoku, where the real world is **identical** to the model.)
- We might formulate the driving problem as a route finding problem as a graph where nodes represent points in the city, edges represent the roads, and the cost of an edge represents the traffic on that road.
- If we do this, all the complexities of perception are **ignored**. Alternatively, we could make a model that only looks at perception and ignores planning.

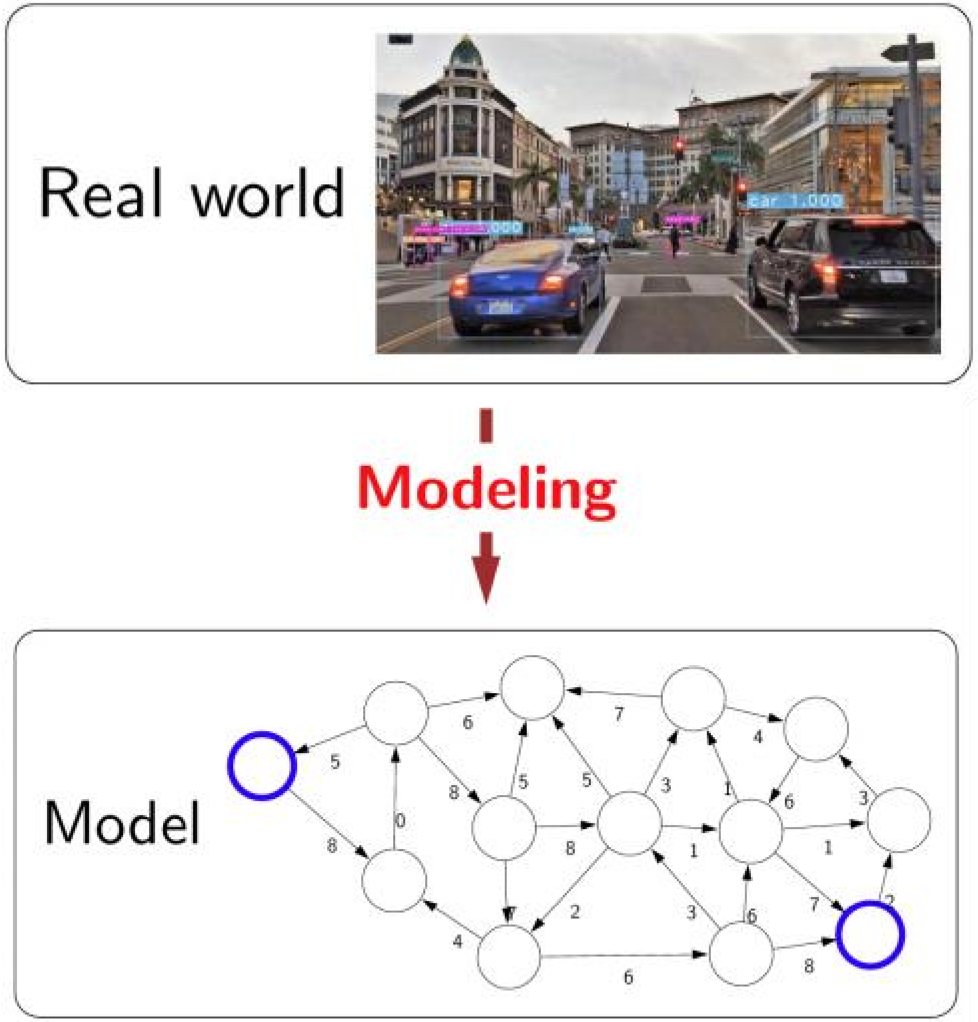

Paradigm: Inference

- The second pillar is **inference**. Given a model, the task of inference is to **answer questions** with respect to the model. For example, given the model of the city, one could ask questions such as: what is the shortest path? what is the cheapest path?
- The focus of inference is usually on **efficient algorithms** that can answer these questions.
- For some models, computational **complexity** can be a concern (games such as Go), and usually approximations are needed

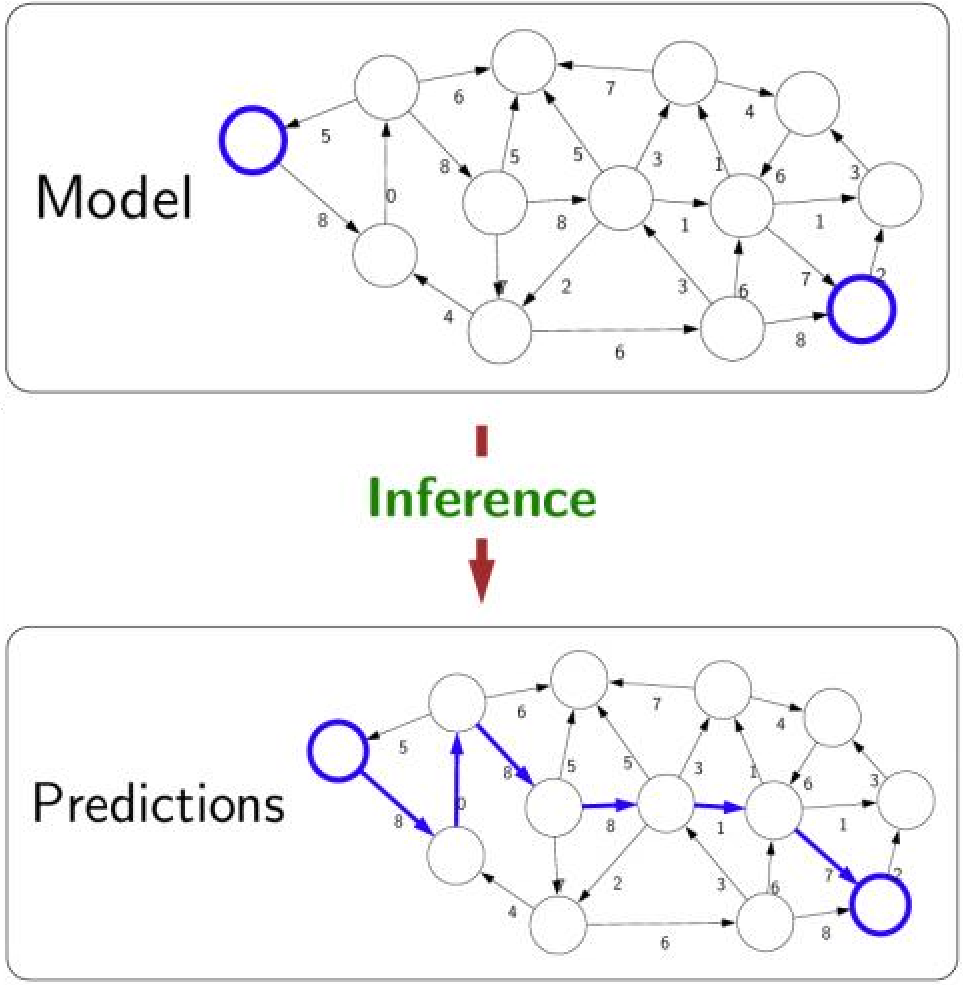

Paradigm: Learning

- But where does the model come from? Remember that the real world is **rich**, so if the model is to be faithful, the model has to be rich as well. But we cant possibly write down such a rich model manually.

- The idea behind (machine) learning is to instead **get it from data**. Instead of constructing a model, one constructs a skeleton of a model (more precisely, a model family), which is a model without parameters. And then if we have the appropriate data, we can run a machine learning algorithm to **tune the parameters** of the model.

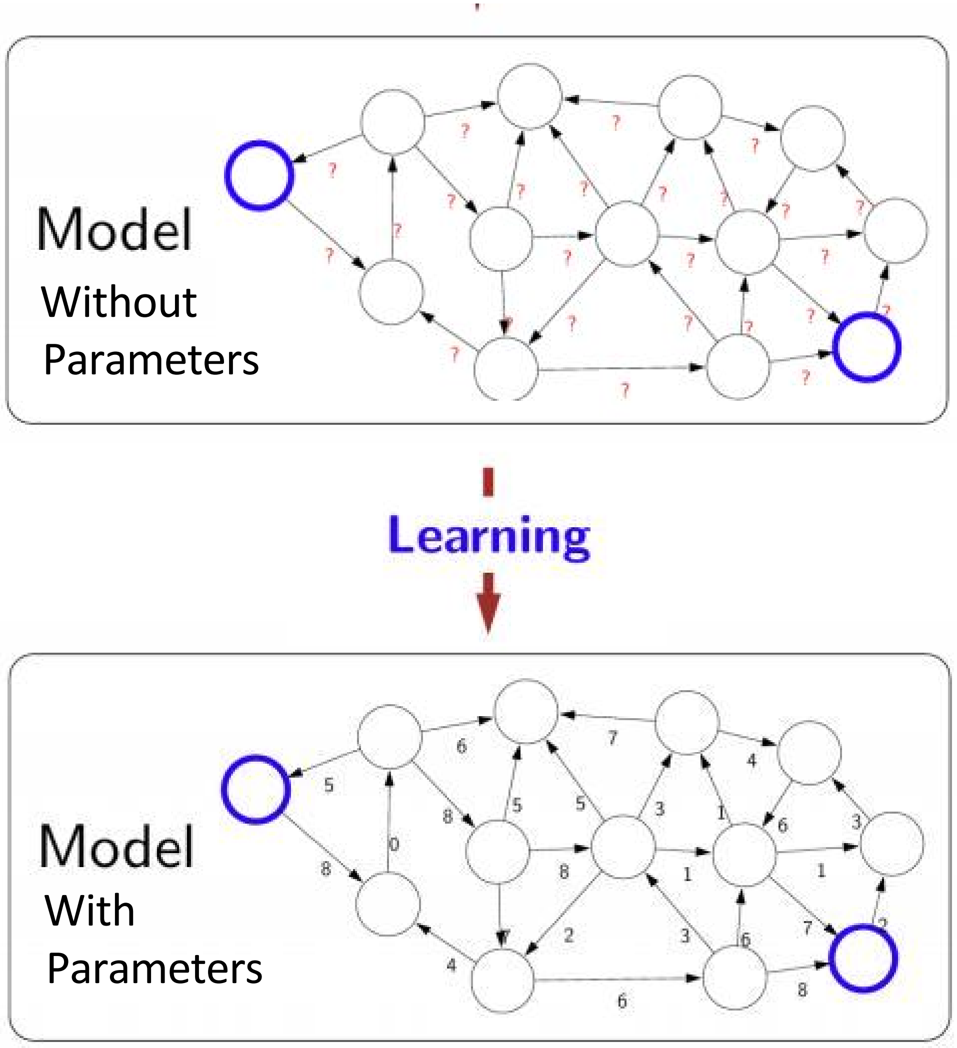

How to Design an AI system？

（1）“What are you attempting to solve for?”
（2）“What is the desired outcome?”

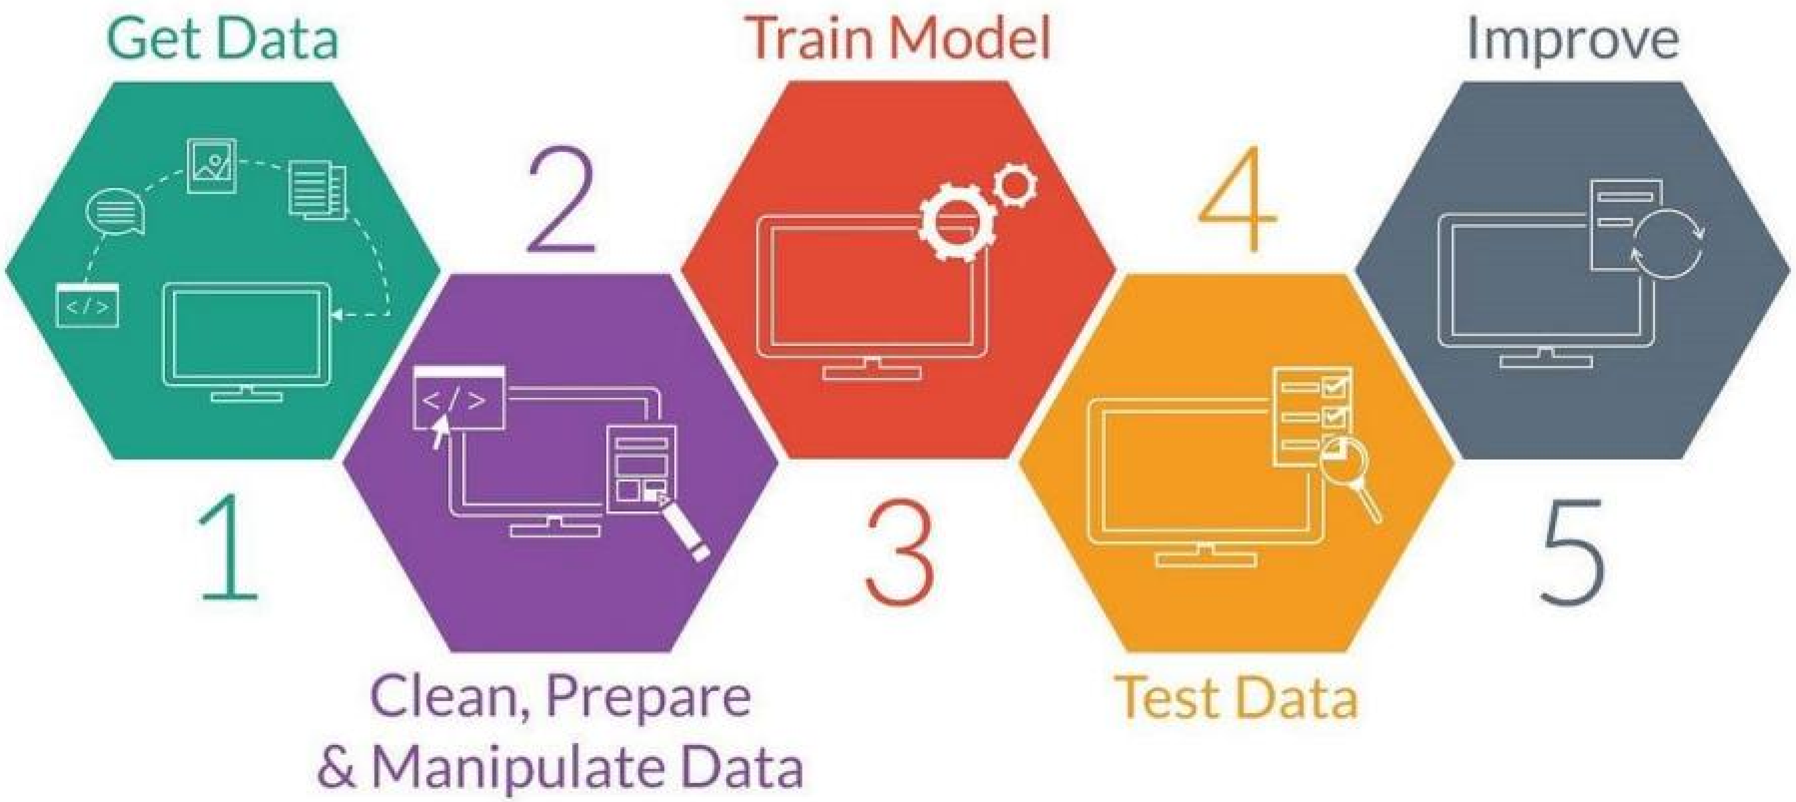

How to Design an AI system?

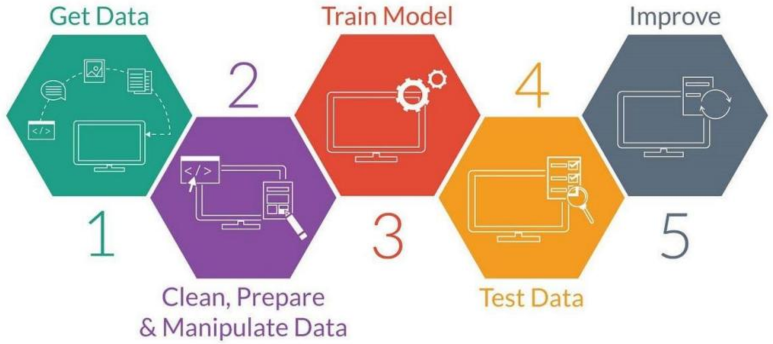

1. **Problem Definition and Requirement Analysis**
   - Identify the problem or goal that the AI system aims to address.
   - Understand the specific requirements, constraints, and objectives of the project
   - Define the scope of the AI system and the expected outcomes.

2. **Data Collection**
   - Determine the data needed to train and evaluate the AI system.
   - Gather relevant and representative data from various sources.
   - Ensure data quality, considering factors like accuracy, completeness, and bias.

3. **Data Cleaning and Preprocessing**
   - Preprocess the collected data to handle missing values, outliers, and inconsistencies.
   - Normalize or standardize the data to ensure compatibility and fairness.
   - Perform data augmentation techniques if necessary to increase the diversity of the dataset.

4. **Data Analysis and Visualization**
   - Explore and analyze the data to gain insights and identify patterns.
   - Utilize statistical methods and visualization techniques to understand the data characteristics.
   - Identify correlations between variables and potential dependencies.

How to Design an AI system?

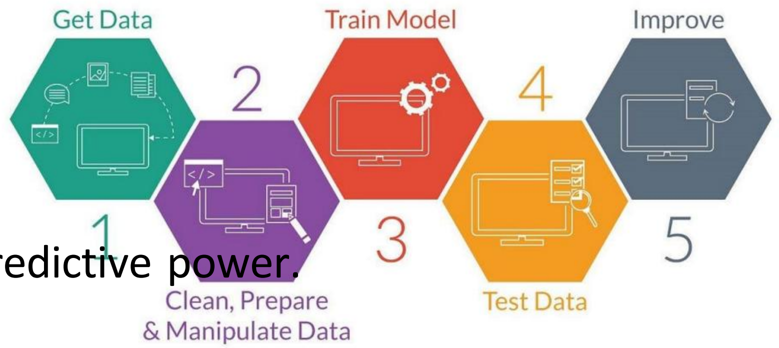

5. **Feature Engineering**
   - Extract or create relevant features from the data that can enhance the model's predictive power.
   - Transform the data to a format suitable for the selected AI model.
   - Apply techniques such as dimensionality reduction, feature selection, or feature scaling.

6. **Model Construction and Training**
   - Select an appropriate AI model or algorithm that aligns with the project requirements.
   - Split the dataset into training, validation, and testing sets.
   - Train the model using the training data and optimize its performance through hyperparameter tuning.

7. **Model Evaluation and Tuning**
   - Validate the model using the validation set and assess its performance metrics.
   - Iterate on the model, adjusting hyperparameters or trying different algorithms as needed.
   - Fine-tune the model to improve its predictive power and generalization capabilities.

8. **Model Results Display and Report Output**
   - Evaluate the trained model using the testing dataset to assess its performance.
   - Generate reports or visualizations to communicate and interpret the results.
   - Present the findings and insights derived from the AI system's predictions or classifications.
   - Provide reports or outputs to stakeholders, decision-makers, or end-users to inform decision-making processes.

How to Design an AI system?

9. **Model Deployment and Online Feedback**
   - Integrate the trained model into a production environment or application.
   - Monitor the performance and behavior of the deployed model in real–world scenarios.
   - Collect feedback from users or systems interacting with the AI system.
   - Continuously improve the model based on feedback and new data to enhance its effectiveness.

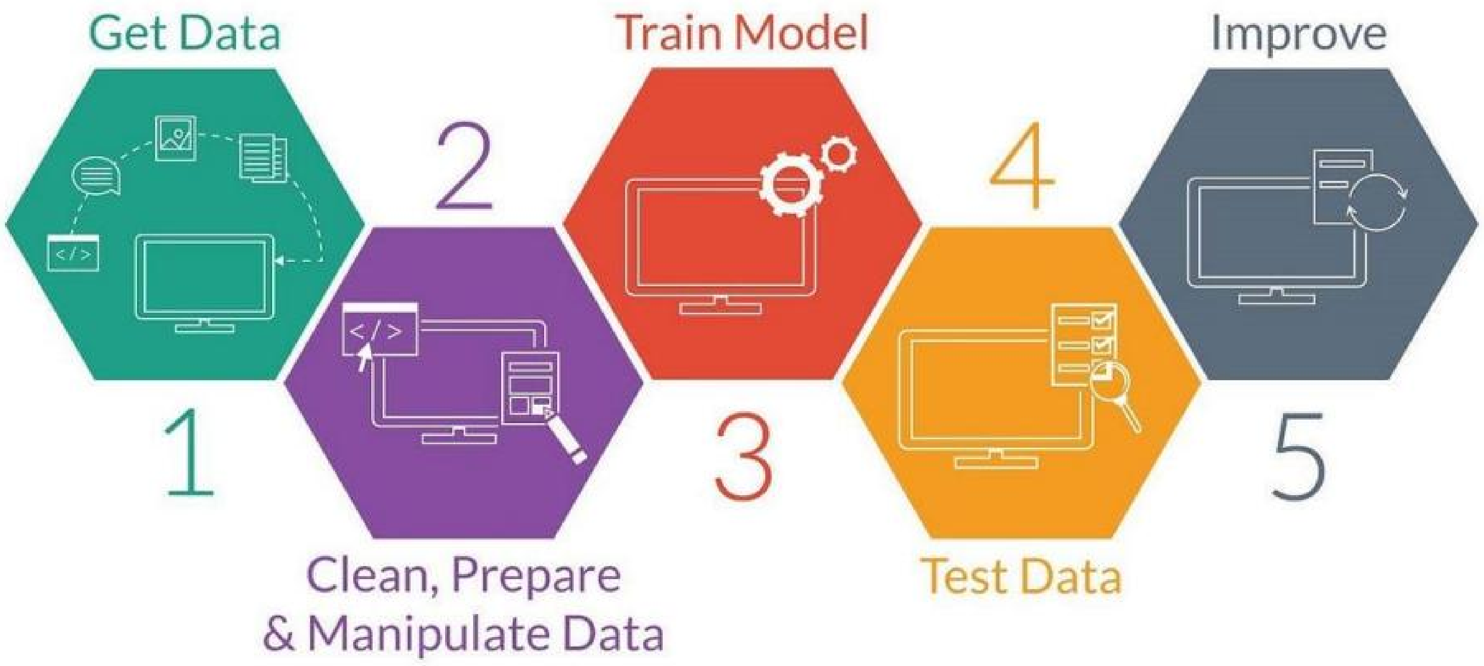

Outline

- AI as Research

AI as Research

AI Research	AI Engineering
Asks “why”	Asks “how”
Seeks general principles	Optimizes specific systems
Publishes theories & models	Deploys products
Long-term impact	Short-term performance

AI as Research

中国计算机学会推荐国际学术期刊

（人工智能）

一、A 类

序号	刊物简称	刊物全称	出版社	网址
1	AI	Artificial Intelligence	Elsevier	http://dblp.uni-trier.de/db/journals/ai/
2	TPAMI	IEEE Transactions on Pattern Analysis and MachineIntelligence	IEEE	http://dblp.uni-trier.de/db/journals/pami/
3	IJCV	International Journal of Computer Vision	Springer	http://dblp.uni-trier.de/db/journals/ijcv/
4	JMLR	Journal of Machine Learning Research	MIT Press	http://dblp.uni-trier.de/db/journals/jmlr/

中国计算机学会推荐国际学术会议

（人工智能）

一、A 类

序号	会议简称	会议全称	出版社	网址
1	AAAI	AAAI Conference on Artificial Intelligence	AAAI	http://dblp.uni-trier.de/db/conf/aaai/
2	NeurIPS	Conference on Neural Information Processing Systems	MIT Press	http://dblp.uni-trier.de/db/conf/nips/
3	ACL	Annual Meeting of the Association for Computational Linguistics	ACL	http://dblp.uni-trier.de/db/conf/acl/
4	CVPR	IEEE/CVF Computer Vision and Pattern Recognition Conference	IEEE	http://dblp.uni-trier.de/db/conf/cvpr/
5	ICCV	International Conference on Computer Vision	IEEE	http://dblp.uni-trier.de/db/conf/iccv/
6	ICML	International Conference on Machine Learning	ACM	http://dblp.uni-trier.de/db/conf/icml/
7	IJCAI	International Joint Conference on Artificial Intelligence	Morgan Kaufmann	http://dblp.uni-trier.de/db/conf/ijcai/

AI as Research

https://www.ijcai.org/

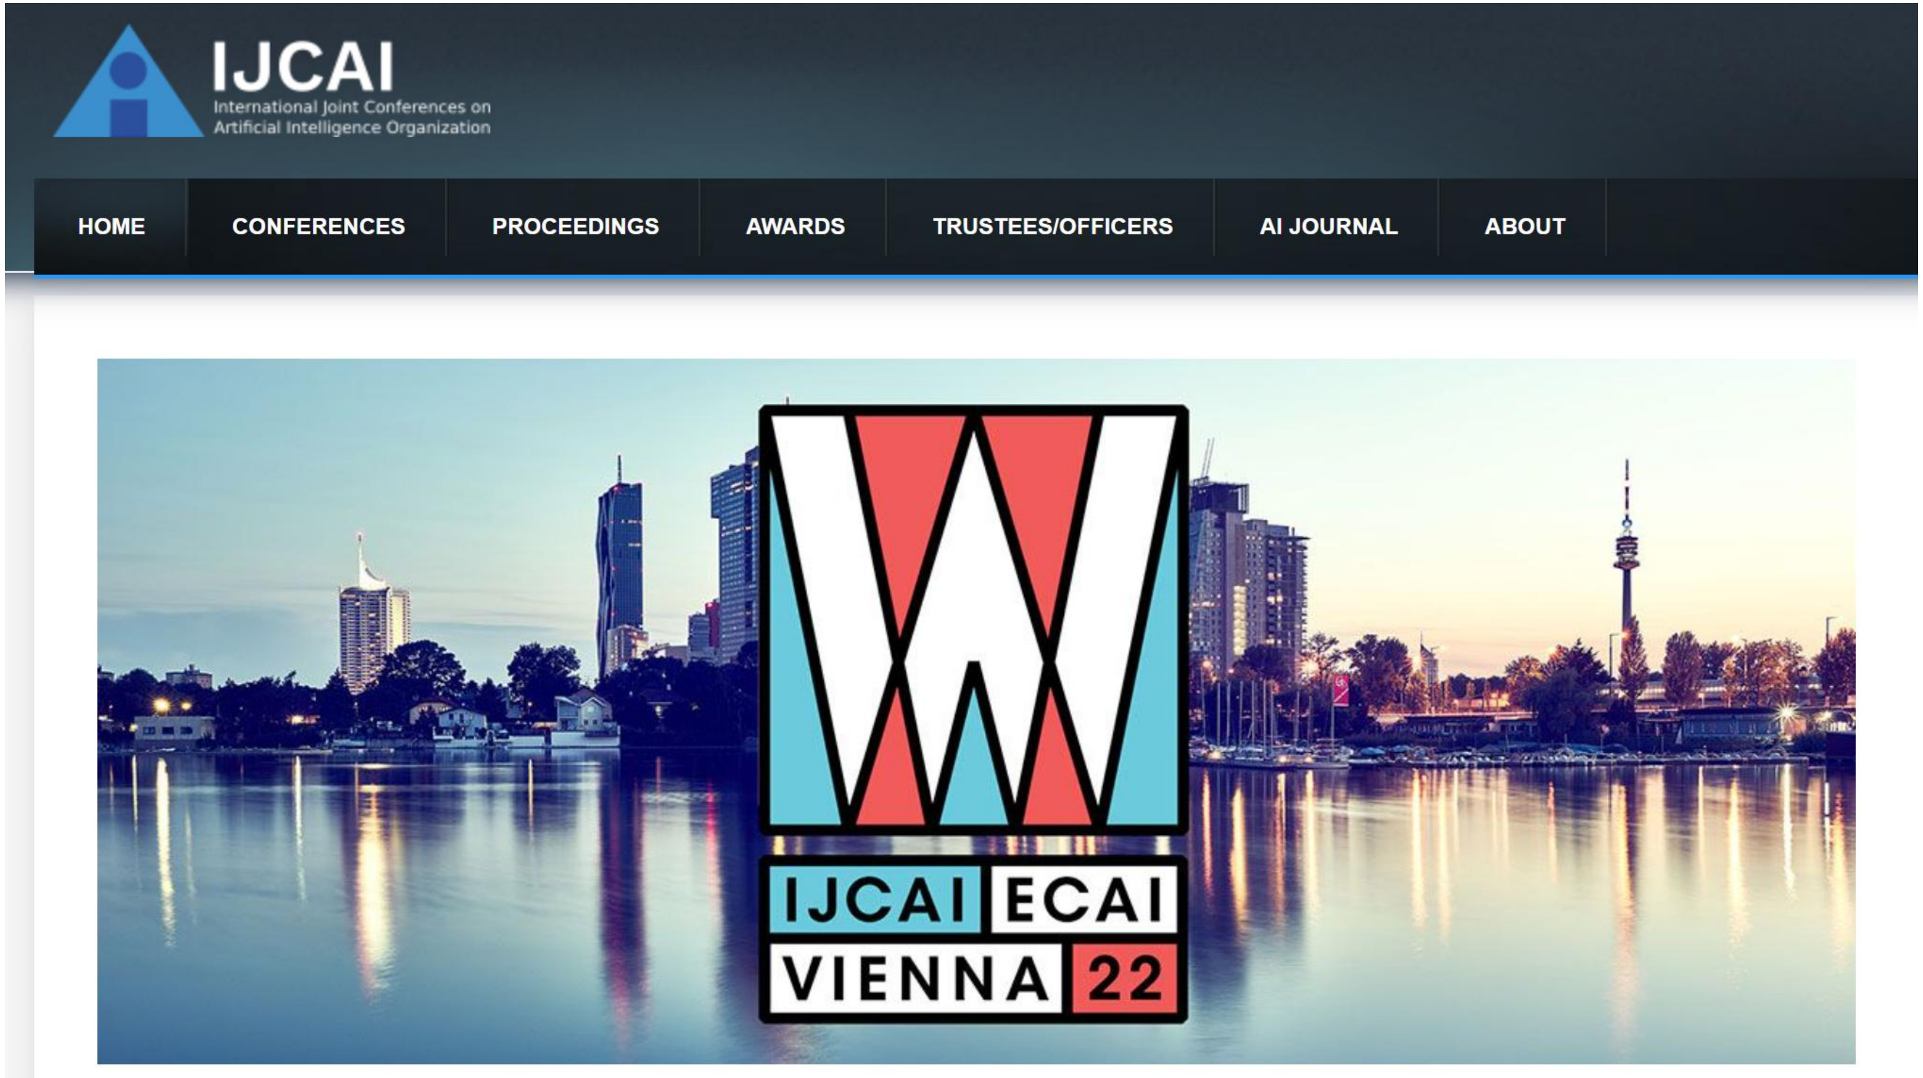

AI as Research

https://aaai.org/conference/aaai/aaai-26/

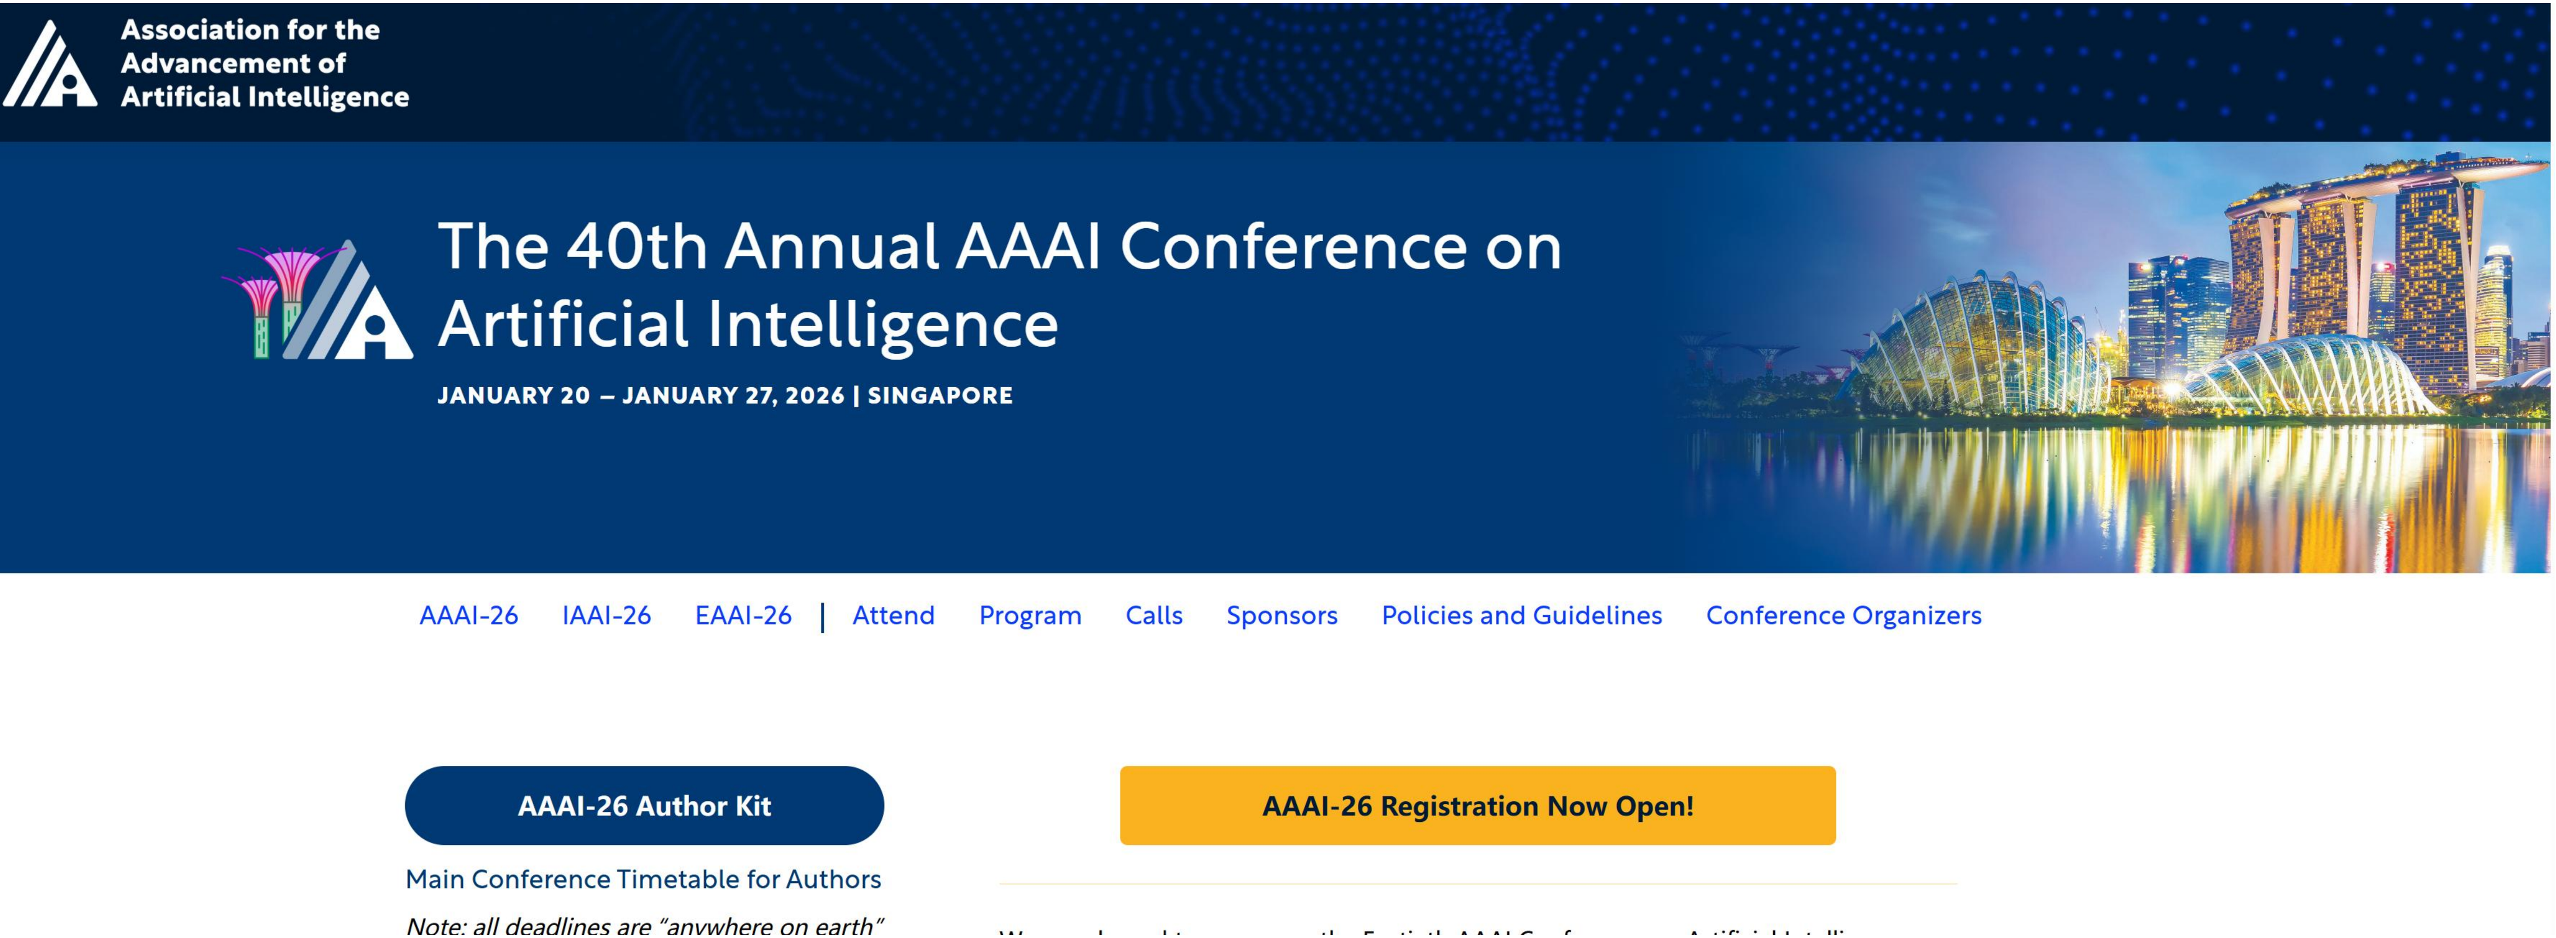

AI as Research

https://iccv.thecvf.com/

International Conference on Computer Vision, ICCV 2025

ICCV is the premier international computer vision event comprising the main conference and several co-located workshops and tutorials.

We expect ICCV 2025 to happen in person at the Hawaii Convention Center in Honolulu.

Hawai'i Convention Center

Oct 19 – 23th, 2025, Honolulu, Hawai'i

Videos

Keynote and Oral videos are available now. Workshop and Tutorial videos (where available) are still being processed. As videos come online, they will be available at the link below.

Videos

Conference Schedule

Tutorials | Workshops | Keynotes | Papers | Oral Sessions

Badge Pickup, Breaks, Reception

Build Personalized Program | Paper Visualizer

AI as Research

https://www.computer.org/csdl/journal/tp

IEEE TRANSACTIONS ON
PATTERN ANALYSIS AND
MACHINE INTELLIGENCE

From the February 2026 Issue: Schedule-Robust Continual Learning

FEATURED ARTICLE

Trending Articles

- Event-Based Vision: A Survey
- Schedule-Robust Continual Learning
- Mask-DiFuser: A Masked Diffusion Model for Unified Unsupervised Image Fusion
- An End-to-End Depth-Based Pipeline for Selfie Image Rectification
- T2I-CompBench++: An Enhanced and Comprehensive Benchmark for Compositional Text-to-Image Generation

Calls for Papers

Announcements

AI as Research

https://scholar.google.com/schhp?hl=en

My profile　My library　Labs

提出详细的研究问题以查找相关论文

问题示例

- Has anyone used single molecule footprinting to examine transcription factor binding in human cells?
- Are hydrogen powered cars, compared to electric / internal combustion engine cars, really better for the environment?
- What is the standard of care for intraductal papilloma without atypia? When is surgical excision recommended, and when can it be managed conservatively?
- Find papers from the past 2 years about how to determine whether an abstractive summary generated by an LLM is grounded.

Google Scholar

Articles　Case law

New! Scholar Labs: An AI Powered Scholar Search

What constitutes a good paper, a good journal, and a good conference?

First is **academic rigor（学术性）**. This refers to research and discussion that is specialized and systematic, focusing on a specific professional issue within a scientific field.

Second is **scientific validity（科学性）**. This characteristic is inherent to the nature of academic papers.

Third is **innovation（创新性）**. A paper lacking innovation holds no value whatsoever.

Fourth is **theoretical depth（理论性）**. The defining distinction between academic papers and popular science works, practice reports, or technical intelligence lies in their theoretical character.

Fifth is **standardization（规范性）**. To facilitate communication and application, papers must employ standardized linguistic systems and symbolic notation for expression.

CCF: 中国计算机学会

中科院期刊分区表

Impact Factor Letpub

Count Citations: Find the number of times articles published in the journal in 2023 and 2024 were cited in 2025.
Count Articles: Find the total number of scholarly articles (citable items) published in the journal in 2023 and 2024.
Divide: Divide the citation count (Step 1) by the article count (Step 2)

AI as Research

The LaTeX Project

Home | About | Get | LaTeX3 | Publications | Help | News

LaTeX – A document preparation system

LaTeX is a high-quality typesetting system; it includes features designed for the production of technical and scientific documentation. LaTeX is the de facto standard for the communication and publication of scientific documents. LaTeX is available as free software.

You don't have to pay for using LaTeX, i.e., there are no license fees, etc. But you are, of course, invited to support the maintenance and development efforts through a donation to the TeX Users Group (choose LaTeX Project contribution) if you are satisfied with LaTeX.

You can also sponsor the work of LaTeX team members through the GitHub sponsor program at the moment for Frank, David and Joseph. Your contribution goes without any reductions by GitHub to the developers in support of the project.

The volunteer efforts that provide you with LaTeX need financial support, so thanks for any contribution you are willing to make.

Recent News

11 November, 2025
LaTeX 2025-11-01 released and distributed

30 October, 2025
Accessibility of STEM documents - talk at PDF days 2025 in Berlin

9 October, 2025
Pre-release 3 of LaTeX 2025-11-01 is available for testing

2 June, 2025
LaTeX 2025-06-01 released and distributed

25 May, 2025
Retirement of the LaTeX-L mailing list

21 May, 2025
Pre-release 4 of LaTeX 2025-06-01 is available for testing

5 May, 2025
Final pre-release of LaTeX 2025-06-01 is available for testing

All News · Subscribe to our RSS News Feed

AI as Research

https://www.overleaf.com/

Product ▾ Solutions ▾ Templates Pricing Help & reso

Back to all gallery items

Filters: **All** / Templates / Examples / Community articles

LaTeX templates and examples — Journal articles

Discover a wide range of academic journal LaTeX templates for articles and papers which automatically format your manuscripts in the style required for submission to that journal.

Recent

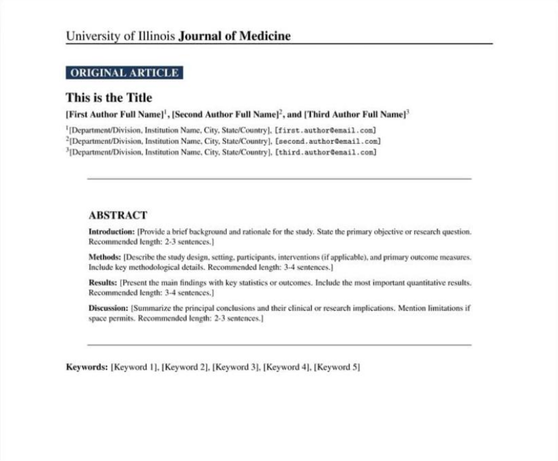

University of Illinois **Journal of Medicine**

ORIGINAL ARTICLE

This is the Title

[First Author Full Name][1], [Second Author Full Name][2], and [Third Author Full Name][3]

[1][Department/Division, Institution Name, City, State/Country], [first.author@email.com]
[2][Department/Division, Institution Name, City, State/Country], [second.author@email.com]
[3][Department/Division, Institution Name, City, State/Country], [third.author@email.com]

ABSTRACT

Introduction: [Provide a brief background and rationale for the study. State the primary objective or research question. Recommended length: 2-3 sentences.]

Methods: [Describe the study design, setting, participants, interventions (if applicable), and primary outcome measures. Include key methodological details. Recommended length: 3-4 sentences.]

Results: [Present the main findings with key statistics or outcomes. Include the most important quantitative results. Recommended length: 3-4 sentences.]

Discussion: [Summarize the principal conclusions and their clinical or research implications. Mention limitations if space permits. Recommended length: 2-3 sentences.]

Keywords: [Keyword 1], [Keyword 2], [Keyword 3], [Keyword 4], [Keyword 5]

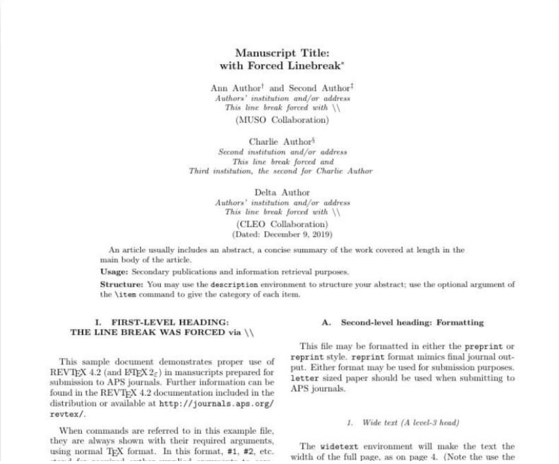

Manuscript Title:
with Forced Linebreak*

Ann Author† and Second Author‡
Authors' institution and/or address
*This line break forced with *
(MUSO Collaboration)

Charlie Author§
Second institution and/or address
This line break forced and
Third institution, the second for Charlie Author

Delta Author
Authors' institution and/or address
*This line break forced with *
(CLEO Collaboration)
(Dated: December 9, 2019)

An article usually includes an abstract, a concise summary of the work covered at length in the main body of the article.

Usage: Secondary publications and information retrieval purposes.

Structure: You may use the `description` environment to structure your abstract; use the optional argument of the `\item` command to give the category of each item.

I. FIRST-LEVEL HEADING:
THE LINE BREAK WAS FORCED via \\

This sample document demonstrates proper use of REVTeX 4.2 (and LaTeX 2_ε) in mansucripts prepared for submission to APS journals. Further information can be found in the REVTeX 4.2 documentation included in the distribution or available at `http://journals.aps.org/revtex/`.

When commands are referred to in this example file, they are always shown with their required arguments, using normal TeX format. In this format, `#1`, `#2`, etc.

A. Second-level heading: Formatting

This file may be formatted in either the `preprint` or `reprint` style. `reprint` format mimics final journal output. Either format may be used for submission purposes. `letter` sized paper should be used when submitting to APS journals.

1. Wide text (A level-3 head)

The `widetext` environment will make the text the width of the full page, as on page 4. (Note the use the

THE CYBER DEFENSE REVIEW
YEAR, VOL. XX, NO. XX, 1-12
https://doi.org/10.55682/cdr/XXXX-XXXX

RESEARCH ARTICLE

Insert The Title of Your Paper Here

First A. Author[*12], Second B. Author[13], Third C. Author[1]

[1]Army Cyber Institute, West Point, NY, USA
[2]Institution Two, City, State, Country
[3]Institution Three, City, Country

Sample text inserted for demonstration. Insert your abstract here, with no distinctive header. The abstract should be between 150 and 200 words. It should not contain bibliographical references. Your abstract must give readers a brief summary of your article. It should be informative and accessible: indicate the general scope of the article and state the main results obtained and conclusions drawn. The abstract must be complete in itself: it must not contain undefined abbreviations and must not refer to any table, figure, reference or equation numbers. The review process for research articles is double-blind: the submitted document should not include author information and should not include acknowledgments, or mentions (e.g., in citations or discussion of related work) that would make the authorship apparent. Upon acceptance, the author and affiliation information will be added to your paper.

Keywords: Insert 3-5 comma delimited keywords, keyword 2, keyword 3, keyword 4

Assignment 1 (Grades account for 10%)

Search and Download at least 5 papers published in top-tier (CCF A) AI venues over the past 5 years within a specific research area;
Systematically read and summarize these works;

Write a concise literature review report (formated with Latex Via Overleaf).

Detailed information will be released via Canvas soon.

Canvas Submission DDL: 01 March 2026

Thank you!

Innovating into the Future