

AI for Mathematics

AI Foundations → Mathematical Research Projects

Ma Jiajun

School of Math, XMU / Dept. of Math, XMUM

April 28 & May 5, 2026

Part I

The Logic and History of AI Research Methods

What Is Artificial Intelligence?

Artificial intelligence is the broad field of building systems that perform tasks associated with intelligence.

Reasoning

Make inferences from given information.

Learning

Improve behavior from data or feedback.

Language

Understand or generate text and speech.

Perception

Interpret images, sound, or sensor data.

Planning

Choose actions to reach a goal.

Problem solving

Search for solutions under constraints.

AI Is Broader Than LLMs

Correct picture

LLMs are one recent and important family of AI systems.

Wrong shortcut

"AI" does not mean only ChatGPT-style systems.

In this course, LLMs matter as one interface to language, code, and tool workflows.

Vocabulary Map

AI: broad field of machine intelligence

Machine learning: behavior learned from data


Deep learning: machine learning with large multi-layer neural networks

Generative AI: systems that generate text, images, audio, code, or other content

LLM: generative AI model specialized in language and code-like token sequences

Historical Timeline

Period	Main idea
1950	Turing's imitation game
1955/1956	Dartmouth proposal/workshop
1950s-1970s	Symbolic AI, search, planning
1980s	Expert systems
1990s-2000s	Statistical machine learning
2010s	Deep learning at scale
2017	Transformer architecture
2020s	LLMs and tool-using agents



Alan Turing's 1950 paper made machine intelligence a concrete scientific question.

Two Long-Running Impulses

Symbolic

Represent knowledge explicitly with rules, logic, search, and plans.

Statistical

Learn behavior from data using probability, optimization, and models.

Modern systems often combine both: neural models plus search, retrieval, code execution, or proof checking.

Method 1: Symbolic AI / GOFAI


Core idea: represent knowledge with symbols, rules, and logic.

Strength

Transparent reasoning; good when rules are explicit and the domain is formal.

Limitation

Brittle when the world is messy, ambiguous, or too large to encode by hand.



Symbolic AI inherited confidence from logic, formal systems, and explicit rules.

Example: Logic Theorist

Background


Logic Theorist was built by Allen Newell, Herbert Simon, and Cliff Shaw in 1956, at the beginning of AI as a named research field.

What it did

It searched for proofs of propositional logic theorems from *Principia Mathematica*. The task was narrow, formal, and rule-governed.

Why it matters

It made the early AI dream concrete: intelligence as explicit symbol manipulation plus heuristic search.



Formal proof was one of AI's earliest test domains.

Method 2: Search

Core idea: formulate a task as states and actions, then explore possible paths until a goal is reached.

Component	Meaning
State	a compact description of where we are
Action	one legal move from a state
Successor function	generates next states
Goal test	decides whether the task is solved
Search strategy	chooses which unfinished node to expand

A search algorithm maintains a frontier of unfinished nodes. The strategy is what makes breadth-first, depth-first, best-first, and A* different.

Tree Search

Node

A node records a state plus the path that reached it. The same state can appear in more than one path unless duplicates are detected.

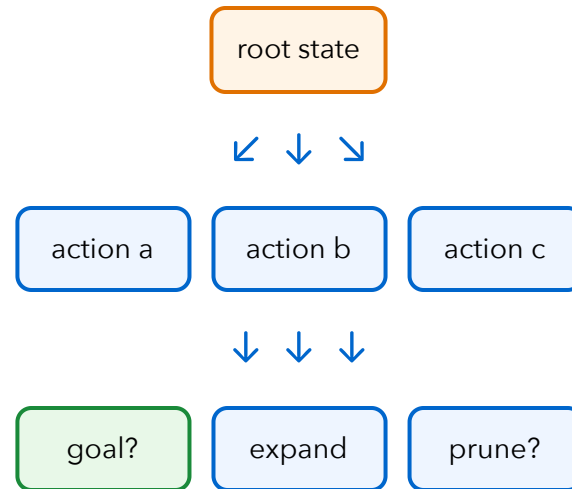
Expansion

To expand a node, apply each legal action and create child nodes.

Frontier

The frontier is the set of generated but not-yet-expanded nodes. Search is mostly the problem of choosing from this frontier.

A search tree



The tree is a bookkeeping structure for possible histories, not necessarily a picture of the whole world.

Proof Search as Tree Search

Search object

In theorem proving, a node can be a proof state: current hypotheses plus goals still to prove.

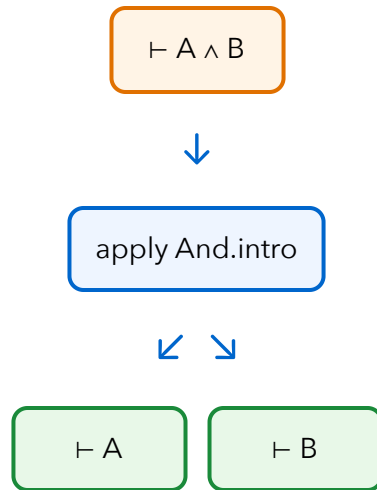
Edges

An inference rule or tactic transforms one goal into zero or more subgoals.

Success condition

A branch succeeds only when every leaf goal is closed by an assumption, axiom, lemma, computation, or decision procedure.

Example proof tree



Lean tactics expose this process interactively; automated provers search much larger proof spaces.

Example: Deep Blue

Background

IBM's Deep Blue defeated world chess champion Garry Kasparov in a 1997 match.

What made it work

It searched enormous chess game trees, evaluated board positions, and used specialized hardware plus chess knowledge.

Method point

Search is powerful when rules are fixed, actions are discrete, and states can be evaluated.



IBM Deep Blue chess computer.



Garry Kasparov, world chess champion.

The match is a classic example of search, evaluation, and domain engineering.

Method 3: Planning

Planning asks for a legal action sequence, not just a plausible next step.

Input

An initial state, a goal condition, and a model of available actions.

Action model

Each action specifies when it is allowed and how it changes the state.

Output

A sequence a_1, \dots, a_T such that every precondition is satisfied and the final state satisfies the goal.

STRIPS-style action

preconditions: facts that must already hold

add list: facts made true by the action

delete list: facts made false by the action

State update:

$$\text{new state} = (\text{old state} - \text{delete}) + \text{add}$$

Planning is search over action sequences plus legality checks at each step.

Example: Shakey and STRIPS

Background

Shakey was an early mobile robot project at SRI in the late 1960s and early 1970s.

What the planner had to do

Represent facts about rooms, doors, the robot, and boxes; then choose actions whose preconditions become true in order.

Why this is planning

A command such as "put the box in the target room" is not one action. It must be decomposed into moves, pushes, and checks.

Toy STRIPS action

At(robot, A)

At(box, A)



Push(box, A, B)



At(robot, B)

At(box, B)

The plan is valid only if every action's preconditions hold at that step.

Method 4: Knowledge-Based / Expert Systems

Core idea: answer questions or make decisions by combining stored knowledge with rules, retrieval, and scoring.

Classic expert systems

Hand-written rules work best in narrow domains with stable concepts and expert heuristics.

Modern knowledge systems

Question-answering systems often mix knowledge bases, search, NLP, and statistical confidence.

Example: IBM Watson on Jeopardy!

Background


In 2011, IBM Watson defeated Jeopardy! champions Ken Jennings and Brad Rutter.

How it worked

It parsed clues, searched large knowledge sources, generated candidate answers, and ranked them by confidence.

System point

This is not a pure rule-based expert system: it shows how knowledge, language processing, retrieval, and statistical scoring can be combined.



IBM Watson prototype hardware. Photo: Clockready / Wikimedia Commons.

Pipeline: clue parsing → evidence retrieval → candidate answers → confidence ranking.

Method 5: Probabilistic AI

Core idea: represent uncertainty explicitly, then update probabilities when evidence arrives.

Hidden variables

The system may care about causes that are not directly observed.

Noisy evidence

Tests, sensors, and reports can be incomplete or wrong.

Posterior belief

Inference computes probabilities after conditioning on observations.

The output is often not "yes" or "no", but a distribution such as $p(\text{cause} \mid \text{evidence})$.

Example: Bayesian Network for Diagnosis

Graph

Nodes are random variables. Arrows show which variables directly affect the local conditional probabilities.

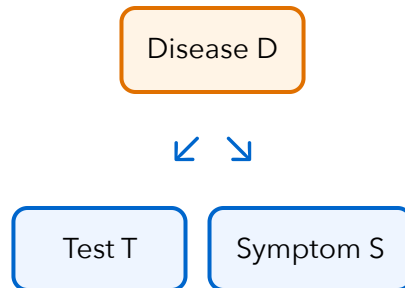
Factorization

For Disease D, Test T, and Symptom S, the graph below represents $p(D,T,S)=p(D)p(T|D)p(S|D)$.

Inference question

After observing T=positive and S=yes, compute $p(D | T=\text{positive}, S=\text{yes})$.

A small Bayesian network



In this simplified model, T and S are conditionally independent given D.

Method 6: Statistical Machine Learning

Core idea: learn patterns from data using statistics and optimization.

examples



fit model



test on unseen data

This shifts the problem from writing all rules to choosing data, models, objectives, and evaluations.

Example: Spam Filters

Background

Email spam filtering was a practical success story for statistical machine learning in the 1990s and 2000s.

How it worked

Systems learned from labeled emails. Words, links, headers, and patterns became features for classifying messages.

Workflow

The workflow is empirical: collect examples, train a classifier, test on new examples, and update as data changes.

Classification from data

labeled emails



classifier




inbox

spam folder

The model is judged by performance on unseen messages.

Method 7: Connectionism and Neural Networks

Core idea: learn distributed representations with many parameters.



Historically powerful but hard to train at scale until better data, hardware, and techniques arrived.

Example: LeNet-5

Background

LeNet-5, developed by Yann LeCun and collaborators in the 1990s, was an influential CNN for handwritten digit recognition.

What it showed

Convolutions reuse small filters across an image, so the model can learn edges, strokes, and digit parts.


Architecture point

CNNs made neural networks practical for visual recognition before the deep learning wave of the 2010s.



Yann LeCun. Photo: Rama / Wikimedia Commons.

CNN / LeNet-Style Architecture



CNNs learn feature maps before classification.

Handwritten Digit Recognition




Historical use case: automated reading of handwritten digits on checks and forms.

LeNet-style systems were built for handwritten digit recognition.

Method 8: Evolutionary Computation

Core idea: search by variation and selection.



Useful for optimization and design search; often sample-inefficient and weak on guarantees.

Example: NASA Evolved Antenna

Background

NASA used evolutionary design methods to produce unusual antenna geometries for space applications.


How it worked

Candidate antenna designs were mutated and selected according to performance criteria. Good designs seeded the next generation.

Search point

Evolutionary computation helps when the design space is strange and human intuition may not suggest the best shape.


Evolutionary design loop



The system searches by keeping variants that perform better.

Method 9: Reinforcement Learning

Core idea: learn actions from reward feedback in an environment.



Hard parts: exploration, reward design, and credit assignment over long action sequences.

Example: AlphaGo

Background

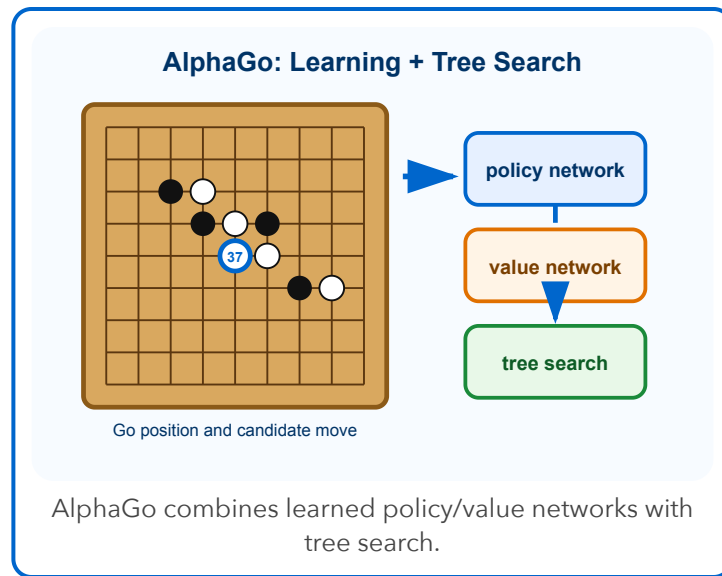
DeepMind's AlphaGo defeated Lee Sedol in 2016, a landmark result because Go has an enormous search space.

What it combined

AlphaGo used deep neural networks, reinforcement learning, and tree search. It learned from games and improved by self-play.

System point

Reinforcement learning is strongest when feedback from actions can become long-term improvement.



Famous context: AlphaGo defeated Lee Sedol 4-1 in Seoul in March 2016.

Method 10: Deep Learning

Core idea: train large multi-layer neural networks on large datasets with gradient-based optimization.

Vision

Image classification,
detection, generation.

Speech

Recognition and synthesis.

Language

Translation, summarization,
generation.

Deep learning can fail silently: confident output is not the same as correct output.

Example: AlexNet

Background

AlexNet won the 2012 ImageNet competition by a large margin and made deep convolutional networks impossible to ignore.


Why it mattered

It benefited from large labeled image datasets, GPUs, and architectural choices that made deep training practical.

Scaling point

Deep learning breakthroughs often come from combining data, compute, architecture, and optimization.


Image classification at scale



The system learns visual features instead of using hand-coded rules.

Method 11: Foundation Models and LLMs

Core idea: pretrain large models on broad corpora, then adapt them to many downstream tasks.



Example: GPT-3 and ChatGPT

Background

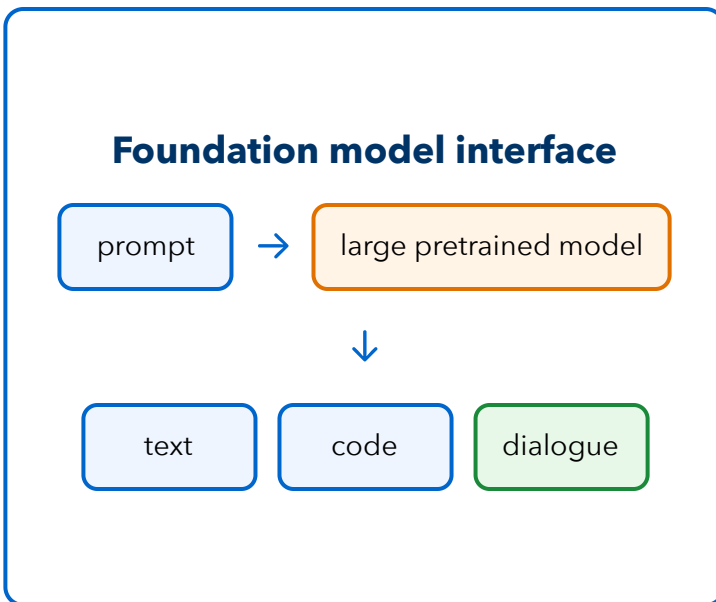
GPT-3 showed that scaling Transformer language models could produce broad few-shot text behavior. OpenAI released ChatGPT publicly in November 2022 as a dialogue interface for instruction-following models.

What changed

The same model interface could draft text, write code, explain concepts, translate, summarize, and interact in natural language.


Interface point

Foundation models are general-purpose interfaces, but fluency is not guaranteed correctness.



Method 12: Agentic and Tool-Using AI

Core idea: put a model inside a loop with tools, memory, retrieval, and evaluation.



Agent systems will appear later when projects need tool use, logs, and external verification.

Example: Claude Code-Style Coding Agent

Background

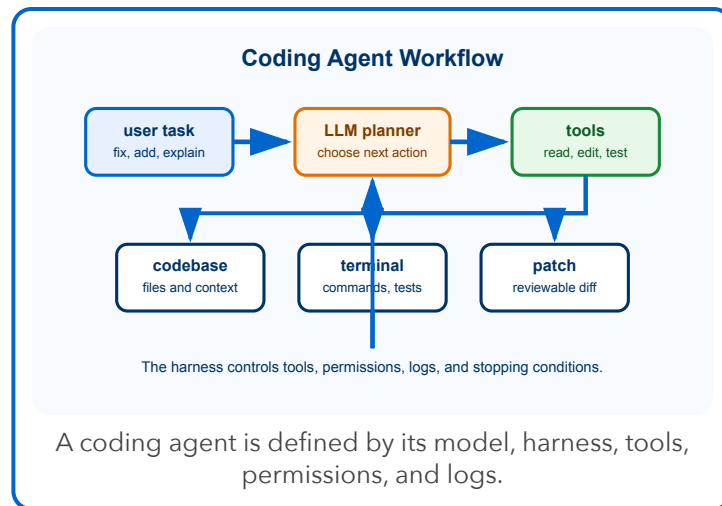
Coding agents such as Claude Code put an LLM inside a development harness that can inspect a repository and use tools.

What changes

The model can read files, search the codebase, edit patches, run commands, inspect errors, and iterate toward a working result.

Agent pattern

ReAct is a classic pattern for reason-plus-action loops; coding agents make the same idea concrete in a real workspace.



Sources: AI History Examples

Claim area	Sources used
Turing and Dartmouth	Turing, " Computing Machinery and Intelligence " (1950); McCarthy et al., " Dartmouth proposal " (1955)
Logic Theorist	Newell, Shaw, and Simon, <i>The Logic Theory Machine</i> (1956/1957); historical account: Gugerty (2006), " Newell and Simon's Logic Theorist "
Shakey / STRIPS	SRI, " Shakey the Robot "; Fikes and Nilsson, " STRIPS " (1971)
Deep Blue	IBM, " Deep Blue "; Campbell, Hoane, and Hsu, " Deep Blue " (2002)
Watson	Ferrucci et al., " Building Watson: An Overview of the DeepQA Project "
Bayesian networks and spam filtering	Pearl, " Probabilistic Reasoning in Intelligent Systems " (1988); Sahami et al., " A Bayesian Approach to Filtering Junk E-Mail " (1998)

Sources: Learning, Games, and Agents

Claim area	Sources used
LeNet-5	LeCun et al., Gradient-Based Learning Applied to Document Recognition (1998)
AlexNet	Krizhevsky, Sutskever, and Hinton, ImageNet Classification with Deep Convolutional Neural Networks (2012)
AlphaGo	Silver et al., Mastering the game of Go with deep neural networks and tree search (2016)
NASA evolved antenna	NASA NTRS, An Evolved Antenna for Deployment on NASA's ST5 Mission
GPT-3 / ChatGPT	Brown et al., Language Models are Few-Shot Learners (2020); OpenAI, Introducing ChatGPT
ReAct / coding agents	Yao et al., ReAct (2022); Anthropic, Claude Code overview

Part II

Language Data, Tokens, and Pretraining

The LLM Training Story

A modern LLM is first trained as a probabilistic sequence model, then often post-trained to behave like an assistant.

Stage	Main object	What is optimized
Tokenization	text -> token IDs	fixed tokenizer chosen before model training
Embedding	token ID -> vector	embedding table learned with the model
Pretraining	Transformer language model	next-token likelihood / cross-entropy
Post-training	assistant behavior	demonstration loss and preference/RL objectives
Decoding	probability -> generated text	inference rule: greedy, sampling, top-p, temperature

What an LLM Learns From

The training data is text. The model learns statistical structure that helps predict continuations in that text.

Published example	Data mixture described in the paper
GPT-3	filtered Common Crawl, WebText2, Books1, Books2, Wikipedia
Preprocessing	quality filtering, deduplication, and tokenization before training
Consequence	learned behavior reflects patterns and errors in the training distribution

Pretraining is not the same as training on a database of verified propositions.

Tokens: The Unit of Text

A token is an element of the model vocabulary. Text is converted to token IDs before entering the neural network.

Word-level problem

A fixed word dictionary cannot handle all misspellings, names, code, and punctuation well.

Character-level problem

Character sequences are long and force the model to relearn word structure.

Tokenization compromise

Use common words, word fragments, punctuation, spaces, and individual characters.

Bishop describes tokenization as a compromise between word-level and character-level representations.

Tokenization Caveats

Language

Different languages may need different numbers of tokens for the same idea.

Symbols

Mathematical notation, code, and rare names may be split into unintuitive pieces.

Context

The context window counts tokens, not characters or pages.

A tokenization problem can become a reasoning problem if important notation is chopped up badly.

Embeddings: IDs Become Vectors

After tokenization, a discrete ID is mapped to a dense vector by an embedding matrix.

If x_n is the one-hot vector for token n , the embedding vector is

$$v_n = Ex_n.$$

Because x_n is one-hot, this is equivalent to selecting one column of E .

Discrete input

Token IDs are symbols, not continuous numbers with natural distance.

Learned table

The columns of E are parameters learned from training data.

Context later

The same token embedding is later changed by Transformer layers using context.

Pretraining Creates Its Own Labels

Self-supervised language modeling turns ordinary text into many prediction tasks.

Token context

Training target

The

capital

The capital

of

The capital of

France

The capital of France

is

No human writes a label for each row. The next token already appears in the text.

Pretraining Objective

For a GPT-style model, the sequence probability is decomposed left-to-right.

$$p_{\theta}(x_1, \dots, x_N) = \prod_{n=1}^N p_{\theta}(x_n \mid x_1, \dots, x_{n-1}).$$

Pretraining chooses parameters that assign high likelihood to observed training sequences:

$$\max_{\theta} \sum_{n=1}^N \log p_{\theta}(x_n \mid x_{<n}).$$

This objective teaches prediction of text continuations. Any facts, grammar, or reasoning patterns are learned only insofar as they help this prediction task.

Loss Function: What Is Minimized?

In practice we minimize negative log-likelihood, usually implemented as cross-entropy loss.

For one token position:

$$\mathcal{L}_t(\theta) = -\log p_\theta(x_t | x_{<t}) = -\sum_{v \in V} y_v \log p_\theta(v | x_{<t}).$$

Vocabulary

All possible token IDs the model can choose from.

Target vector

One-hot target: 1 only for the correct token.

Model distribution

The model's predicted probability over the vocabulary.

Loss Function: A Concrete Example

Context:

The capital of France is

Suppose the correct next token is " Paris".

Good prediction

$p(\text{" Paris"}) = 0.80$ gives loss about 0.22.

Bad prediction

$p(\text{" Paris"}) = 0.01$ gives loss about 4.61.

Training pressure

The optimizer changes weights to make the correct token less surprising.

So the immediate training goal is not "understand France"; it is "assign higher probability to the observed next token in context."

What the Loss Pushes the Model To Do

Raise correct probability

If the real next token is "the", the loss decreases when the model assigns "the" more probability.

Use context

The model must use earlier tokens to predict which continuation is likely.

Average over scale

The final loss averages over many positions, documents, and batches.

This objective does not explicitly say "learn grammar" or "learn facts"; those abilities emerge because they help predict text.

The loss rewards probable continuation, not guaranteed truth. This is why verification is still needed.

The Training Loop

A language model is a parameterized function. Training adjusts its weights to reduce average cross-entropy loss.

Step	Operation
1	sample a batch of token sequences
2	compute next-token probability distributions
3	compute cross-entropy against the real next tokens
4	backpropagate gradients through the Transformer
5	optimizer updates weights

This is statistical learning from text examples, not symbolic proof search.

What Pretraining Gives You

Language

Grammar, style, translation patterns, summarization patterns.

Knowledge

Facts and associations reflected in the training distribution.

Code

Syntax, APIs, conventions, and common debugging patterns.

Reasoning patterns

Examples of explanations, proofs, calculations, and plans.

Transfer

The same model can be prompted for many tasks.

Weakness

Pretraining alone does not guarantee truth, safety, or instruction following.

Part III

Transformer Internals and LLM Training Stages

Transformer: The Core Problem

The same token embedding must become different contextual representations in different sentences.

Input

A sequence matrix X with N rows, one row per token position.

Attention stage

Mix information across rows: which other positions matter for this position?

MLP stage

Transform features inside each row using a shared feed-forward network.

Bishop's description is useful pedagogically: attention mixes token positions; the feed-forward network changes features within each token vector.

Original Transformer Architecture

Encoder side

The original Transformer used stacked encoder layers for the input sequence.

Decoder side

The decoder used masked self-attention plus attention over encoder outputs.

LLM connection

GPT-style language models keep the left-to-right decoder idea and use causal masking.

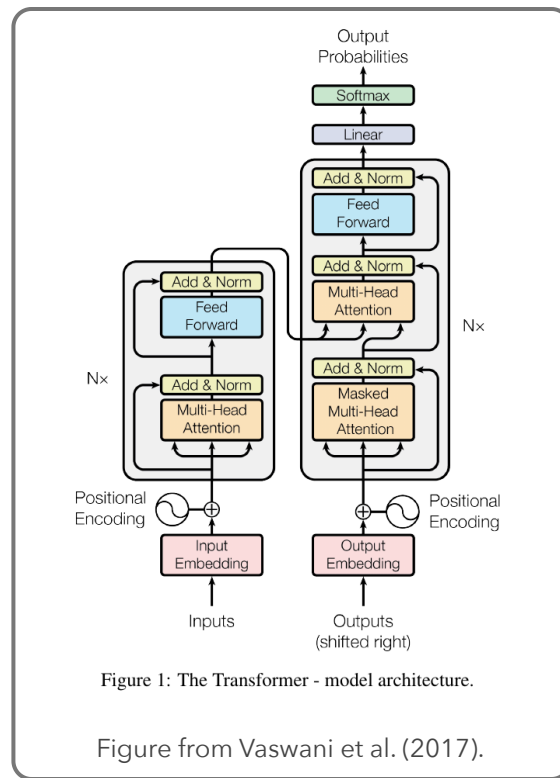


Figure 1: The Transformer - model architecture.

Figure from Vaswani et al. (2017).

Sequence as a Matrix

After embedding, one passage of text is represented by a matrix X .

$$X \in \mathbb{R}^{N \times D}$$

N = number of token positions, D = embedding / model dimension.

Each row x_n^T is the vector representation of one token position.

A Transformer layer maps X to another $N \times D$ matrix. It preserves the number of token positions while changing their representations.

What One Transformer Layer Does

Mix positions

Self-attention lets one token use information from other tokens.

Transform features

A shared MLP applies a nonlinear transformation to each token vector.

Train deeply

Residual connections and layer normalization keep stacked layers stable.

Bishop & Bishop, Chapter 12, describes the same structural split: attention mixes the rows of the sequence matrix; the feed-forward network changes the features inside each row.

So a Transformer block is not only attention. It is attention plus MLP plus normalization/residual structure.

Attention as a Weighted Sum

For output position n , attention forms a new vector by a weighted sum:

$$y_n = \sum_{m=1}^N a_{nm} x_m, \quad a_{nm} \geq 0, \quad \sum_m a_{nm} = 1.$$

Large weight

More information from token m flows into output n .

Small weight

Token m has little influence on output n .

Data-dependent

The weights are computed from the current input, not fixed by hand.

Q/K/V Intuition From Recommendations

Bishop introduces Q/K/V using an information-retrieval analogy: search for a movie using a query and item keys.

Query

User profile or current request: what am I looking for?

Key

Movie/product/post descriptor: what does this item offer?

Value

The content returned or mixed in after matching.

Netflix, Amazon, or Douban are useful analogies, but Transformer Q/K/V are learned vectors inside the network.

Q/K/V Inside a Transformer

Starting from the sequence matrix X , the layer learns three projections:

$$Q = XW_Q, \quad K = XW_K, \quad V = XW_V.$$

QK^T

Every query position is compared with every key position.

softmax rows

Each row becomes attention weights that sum to one.

multiply by V

The weights combine value vectors into contextual outputs.

Scaled Dot-Product Attention

The standard computation used in the Transformer paper is:

$$A = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right), \quad Y = AV.$$

Scores

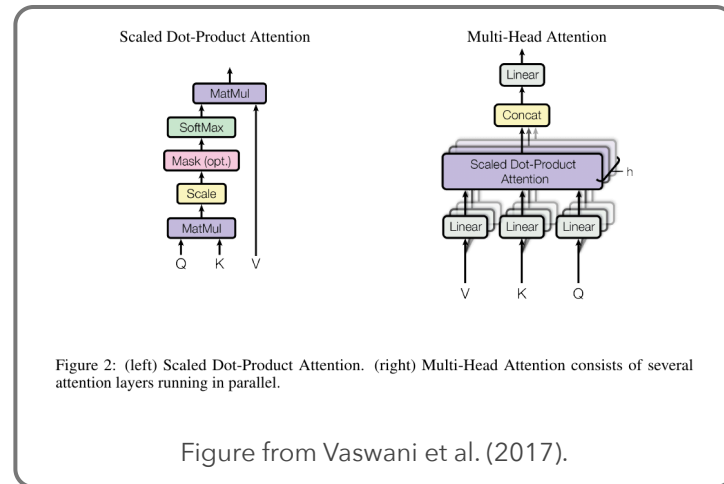
QK^T compares every query position with every key position.

Softmax

Each row becomes an attention distribution over positions.

Output

Multiplying by V forms one updated vector for each position.



Why Divide by $\sqrt{d_k}$?

Dot products grow

If vector components are roughly unit scale, $q \cdot k$ has variance proportional to d_k .

Softmax saturates

Very large positive or negative scores make probabilities close to 0 or 1.

Training suffers

Saturated softmax gives weak gradients and unstable learning.

The factor $\sqrt{d_k}$ keeps attention scores at a reasonable scale before softmax.

This is a numerical design choice that helps optimization; it is not a separate reasoning module.

Multi-Head Attention

Several heads

Each head has its own learned Q , K , and V projections.

Parallel attention

Heads compute attention maps in parallel on the same sequence.

Concatenate + project

Head outputs are concatenated and linearly projected back to dimension D .

Do not overinterpret a head as a human-labeled module. Heads are learned components optimized by the training loss.

The Full Transformer Block

Multi-head attention

Mixes information across token positions.

Residual + normalization

Adds the old representation back and stabilizes scale.

Position-wise MLP

The same feed-forward network transforms each row independently.

Vaswani et al. use sublayer output followed by residual connection and layer normalization. Bishop notes modern variants may use pre-norm, where normalization is applied before the sublayer.

Positional Information

Problem

Self-attention by itself does not encode the order of token positions.

Language needs order

"not good" and "good not" contain similar tokens but different structure.

Solution

Add positional encodings or learned position embeddings to token embeddings.

Transformer Families

Family	Attention pattern	Common use
Encoder	can attend bidirectionally within the input	representation, classification, retrieval
Decoder	causal self-attention: left-to-right generation	GPT-style language modeling
Encoder-decoder	decoder attends to encoder outputs	sequence-to-sequence translation/summarization

GPT-Style Decoder: Training Target

A decoder language model represents a probability distribution for the next token:

$$p(x_1, \dots, x_N) = \prod_{n=1}^N p(x_n \mid x_1, \dots, x_{n-1}).$$

Input

Previous tokens in the context.

Output

A probability distribution over the vocabulary.

Loss

Cross-entropy penalizes low probability on the true next token.

During training the full sequence is known, but the model is masked so it cannot use future tokens to predict the present token.

Causal Mask

Autoregressive rule

Position n may use tokens $1, \dots, n$, but not tokens to its right.

Training


The full sequence is known, but future positions are hidden from each prediction.

Implementation

Future attention scores are set to a very negative value before softmax.

For positions i and j in a decoder layer, the mask permits attention only when $j \leq i$.

Inference: Text Is Generated One Token at a Time



This is autoregressive generation. The model repeatedly extends its own context.

Training uses next-token prediction over many text positions. Inference uses the trained model to repeatedly sample or choose the next token.

Decoding Choices

Greedy

Always choose the most likely next token. Deterministic, often repetitive.

Top-k

Sample only from the k most likely next tokens.

Top-p

Sample from the smallest set whose total probability reaches p.

Decoding is not training. It is the rule used at inference time to turn probabilities into actual text.

Temperature is important enough to discuss separately: it changes the probability distribution before sampling.

Temperature

Before choosing the next token, the model produces logits: one score for each token in the vocabulary.

Temperature rescales these logits before softmax:

$$p_T(v | c) = \frac{\exp(z_v/T)}{\sum_{u \in V} \exp(z_u/T)}.$$

$$T = 1$$

Use the model's ordinary softmax distribution.

$$0 < T < 1$$

Sharper distribution: high-probability tokens dominate more.

$$T > 1$$

Flatter distribution: lower-probability tokens get more chance.

Temperature: Concrete Effect

Suppose the context is:

The proof follows by

and the model's next-token logits favor three continuations.

Token	T = 0.5	T = 1	T = 2
induction	0.879	0.705	0.547
contradiction	0.119	0.260	0.331
compactness	0.002	0.035	0.122

Temperature does not change the model's weights, knowledge, or reasoning ability. It only changes how the next token is sampled.

LLM Training Stages

Pretraining

Question: what continuation is statistically likely?

SFT

Question: how should an assistant answer an instruction?

RL / preference tuning

Question: which answer is preferred under a feedback signal?

These are different stages. Pretraining learns broad language and world patterns; post-training shapes behavior.

None of these stages is the same thing as formal verification.

From Base Model to Assistant

Base model

Pretrained mainly to predict the next token on large text corpora.

SFT

Supervised fine-tuning teaches the model to imitate high-quality instruction responses.

Preference tuning

RLHF-style methods use preferences to push outputs toward desired behavior.

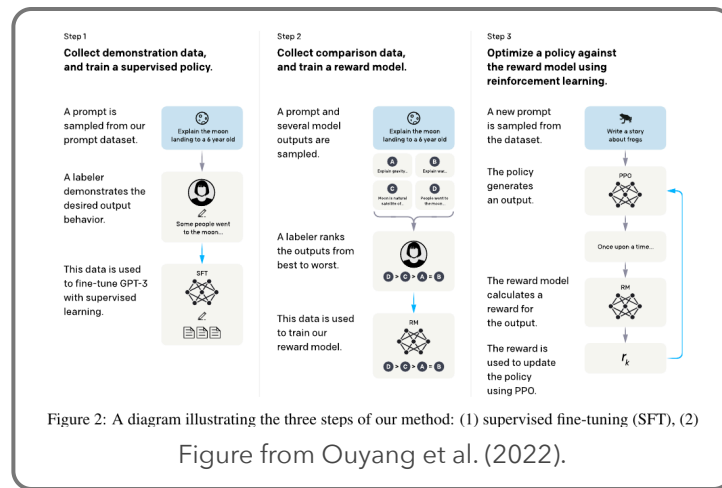


Figure 2: A diagram illustrating the three steps of our method: (1) supervised fine-tuning (SFT), (2)

Figure from Ouyang et al. (2022).

Supervised Fine-Tuning

SFT trains on curated examples: instruction -> good response.

Data

Prompts, tasks, conversations, tool-use examples, and desired answers.

Objective

Usually still next-token loss, but on the desired assistant response.

Effect

Better instruction following, answer format, and task completion habits.

SFT is supervised learning: imitate demonstrations selected by humans or by a data pipeline.

SFT can improve usefulness, but it does not by itself guarantee truth or mathematical correctness.

RL and Preference Tuning in LLMs

Preference data

Two or more model answers are compared; one is marked better.

Reward signal

A reward model or direct preference objective estimates preferred behavior.

Policy update

The model is adjusted to produce preferred outputs while staying close to the base policy.

RLHF often means: collect preferences, train a reward model, then optimize the language model against that reward.

Here "reinforcement learning" means optimizing behavior against a feedback signal. It does not mean the model has discovered mathematical truth.

Keep These Concepts Separate

Transformer

The neural architecture:
attention, MLP, residuals,
normalization, position
information.

Training

How weights are learned:
pretraining, SFT,
preference/RL methods.

Decoding

How text is generated at
inference time: greedy,
sampling, top-p, temperature.

This separation avoids a common confusion: changing temperature is not fine-tuning, and SFT/RL do not change the Transformer block formula.

Why Fluency Is Not Truth

A generated answer can be grammatical, confident, and false.

Objective

Pretraining rewards plausible continuations.

Training data

Contains errors, gaps, and mixed-quality sources.

Post-training

Improves assistant behavior, not mathematical certainty.

Sources: Transformer and Post-Training

Topic	Sources used
Transformer architecture	Vaswani et al., Attention Is All You Need (2017); Bishop and Bishop, <i>Deep Learning: Foundations and Concepts</i> , Ch. 12
Self-attention and multi-head attention	Vaswani et al. (2017), Sec. 3.2; Bishop and Bishop, Ch. 12.1
Next-token language modeling	Brown et al., Language Models are Few-Shot Learners (2020)
SFT and RLHF	Ouyang et al., Training language models to follow instructions with human feedback (2022)
Direct preference objectives	Rafailov et al., Direct Preference Optimization (2023)
Temperature and decoding	Holtzman et al., The Curious Case of Neural Text Degeneration (2019)


Mathematical Discovery and Course Projects

Examples, conjectures, proof critique, formalization, and
verification

Part IV

How Mathematical Discovery Often Happens

A Clean Discovery Pattern



This is the clean version. Real discovery loops backward many times.

What Actually Happens

Compute

Make tables, test small cases, draw diagrams.

Compare

Look for similar behavior and exceptional cases.

Approximate

Find the right scale before the exact law.

Rephrase

Change notation or move to another domain.

Prove

Turn the pattern into necessary implications.

Refine

Fix assumptions and sharpen the statement.

Project-Scale Discovery Tasks

Each item below can become a bounded experiment or tool workflow.

Examples

Generate or compute data.

Patterns

Suggest conjectures.

Reformulation

Find equivalent statements or useful language.

Proof critique

Find gaps and hidden assumptions.

Formalization

Translate into Lean.

Verification

Check claims with tools.

Part V

Prime Number Theorem: A Discovery Case Study

Scope of the PNT Case Study

You do not need analytic number theory today.

Data

Prime tables suggested a pattern.

Conjecture

The pattern became a precise asymptotic statement.

Reformulation

The proof required new analytic language.

We will track numerical evidence, conjectural formulas, and analytic reformulation. The proof itself is outside today's scope.

The Question

Let $\pi(x)$ be the number of primes less than or equal to x .

**How fast does $\pi(x)$
grow?**

Before the theorem was proved, prime tables suggested a stable asymptotic pattern.



Gauss later reported noticing regularity in prime tables long before the theorem was proved.

Count the Primes

x	$\pi(x)$	$x / \log x$
10	4	4.34
100	25	21.71
1,000	168	144.76
10,000	1,229	1,085.74
100,000	9,592	8,685.89

The approximation is not exact, but it has the right order of growth.

Gauss and Legendre

Gauss later wrote to Encke that his work with prime tables in 1792/1793 led him toward logarithmic estimates for prime counting.

Legendre published an approximation for prime counts in 1798 and refined it in later editions.

Accuracy point: they did not prove the Prime Number Theorem. They observed and formulated the asymptotic pattern.



The history starts from observation and approximation, not from a completed proof.

The Emerging Law

$$\pi(x) \sim x / \log x$$

This means that the ratio between $\pi(x)$ and $x / \log x$ tends to 1 as x grows.

Euler's Product

Euler connected primes to an analytic object:


$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s} = \prod_{p \text{ prime}} \frac{1}{1 - p^{-s}}.$$

For this product formula, take $\operatorname{Re}(s) > 1$.

The product says that information about primes is encoded in the zeta function.

Riemann's Reformulation

In 1859, Riemann linked prime counting to complex analysis and the zeros of $\zeta(s)$.



Proof in 1896

Hadamard and de la Vallee Poussin independently proved the Prime Number Theorem in 1896.

data tables

approximate law

analytic reformulation

new proof machinery

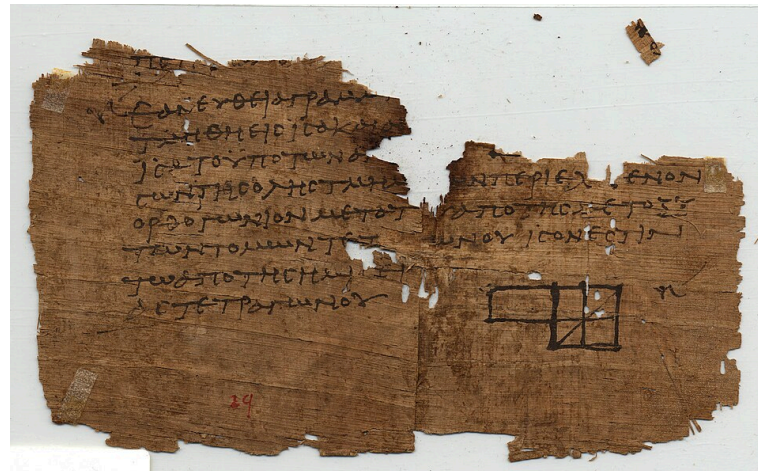
theorem

Project-Scale Tasks in the PNT Story

Do not make the project "prove PNT."

A project-scale version can focus on one artifact: tables, conjectures, reformulations, proof checks, or citation maps.

The object of study must be small enough to evaluate honestly.



The PNT story involves tables, formulas, analytic reformulation, and proof.

Sources: Prime Number Theorem Case

Claim area	Sources used
Gauss, Legendre, Riemann, Hadamard, de la Vallee Poussin	Apostol, The Story of the Prime Number Theorem
Riemann's 1859 memoir	Riemann, On the Number of Prime Numbers less than a Given Quantity
Prime number theorem statement	Encyclopaedia Britannica, Prime number theorem
Euler product formula	Apostol, <i>Introduction to Analytic Number Theory</i> , Ch. 11; standard statement for $\text{Re}(s) > 1$

Part VI

Agents, Harnesses, Tools, Memory, Skills

Why LLM Alone Is Not Enough

A one-shot answer has no reliable feedback loop.

No verifier

The model may sound confident but be wrong.

No search log

We cannot inspect failed attempts.

No tools

It may need computation, retrieval, or Lean.

Model vs Agent

Model


Maps context to output tokens.

prompt → completion

Agent

A model inside a loop that can plan, call tools, observe results, revise, and stop.

The Agent Loop



The observation can be a search result, a computation, a Lean error, or a human rubric.

What Is a Harness?

The harness is the outer program that runs the model or agent.

Instructions

System prompt and policies.

Tools

What the agent can call.

Access

Files, network, and permissions.

Memory

Notes, logs, prior outputs.

Retries

When and how to try again.

Evaluation

How success is judged.

Tools for Math Agents

Search

Papers, web, theorem databases.

Python

Examples, data, counterexamples.

Lean

Formal proof checking.

CAS


Symbolic algebra and simplification.

Retriever

Bring relevant facts into context.

Evaluator


Score outputs on a benchmark.



The four-color theorem is a familiar example where proof, computation, and checking meet.

Retrieval and RAG

Retrieval brings external information into the current context.



RAG means retrieval-augmented generation.

Memory, Skills, Benchmarks, Verifiers

Memory

Stored notes or logs across steps.

Skill

Reusable packaged procedure for a recurring task.

Benchmark

Fixed task set for evaluation.

Verifier

Independent checker such as Lean or a rubric.

Trace

Record of attempts, tool calls, and failures.

Metric

Numerical or structured success criterion.

Example Skill: Proof Review

Input: an informal proof draft.

- Check definitions and quantifiers.
- Identify hidden assumptions.
- Test each implication in order.
- Look for counterexamples to vague claims.
- Separate critical errors from repairable gaps.
- Suggest formal statements that could verify key lemmas.

Part VII

Course Project Directions

Project Formats

Benchmark study

Benchmark and analysis: prompts, examples, metrics, error taxonomy.

Tool workflow

Small tool workflow: retrieval, checker, or retry loop.

Agent prototype

Agent prototype: planning, tools, logs, evaluation, and failure analysis.

The format is flexible. The key requirement is a clear question, a reproducible setup, and honest evaluation.

What Makes a Good Project?

A good project is small, measurable, and reproducible.

Question

One sentence, testable.

Dataset

Concrete examples or problems.

Baseline

A simple comparison method.

Metric

How success and failure are measured.

Artifact

Code, prompts, benchmark, report, demo.

Failure analysis

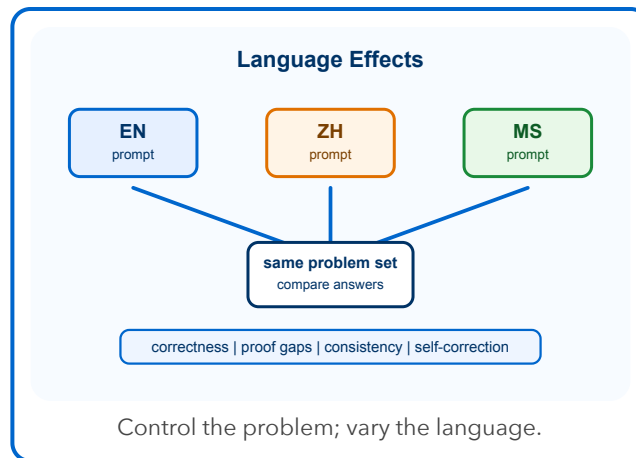
What goes wrong, and why?

Project 1: Language Effects

Question: Does prompt or solution language affect mathematical reasoning quality?

- Compare English, Chinese, Malay, and bilingual prompts.
- Use the same problem set and same evaluation rubric.
- Separate translation errors from reasoning errors.

Metrics: correctness, proof gaps, consistency across runs, self-correction rate.



Project 2: Prover Agent Design

Question: Can an agent improve proof generation by planning, checking, and revising?

informal proof strategy


lemma retrieval or theorem search

Lean statement or proof skeleton

verification error and revision log

Metrics: proof success rate, useful lemma rate, number of revisions, human repair time.

Minimal Prover Agent



Measure success rate, attempts, and common Lean errors.

A prover agent separates strategy, retrieval, formalization, checking, and revision.

Rethlas / Archon as Inspiration

This is inspiration for the prover-agent project, not a requirement to reproduce a research system.

Rethlas

Informal reasoning agent for exploring strategies.

Matlas

Theorem search used by Rethlas.

Archon

Formal verification agent that builds Lean 4 projects.

Possible metrics: useful lemma rate, skeleton quality, type-checking rate, human repair time.

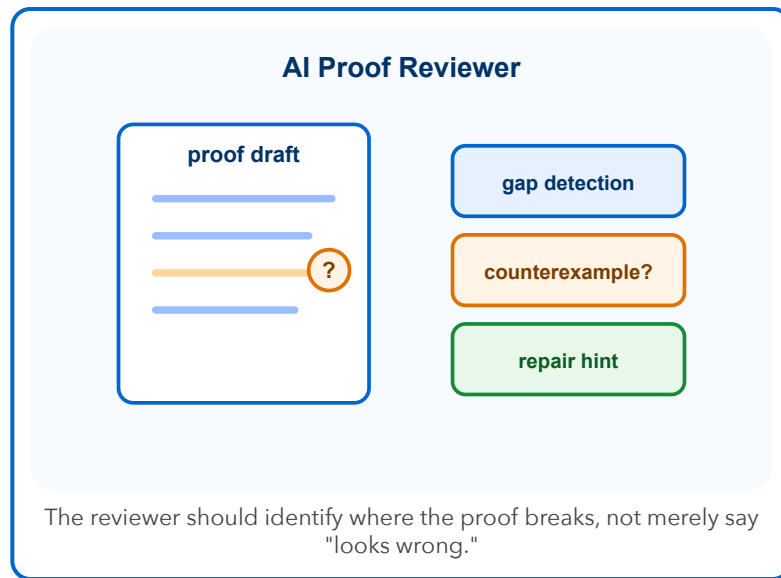
Project 3: AI Proof Reviewer

Question: Can an LLM find gaps in informal mathematical proofs?

Dataset

- Correct proofs.
- Subtly flawed proofs with planted errors.
- Incomplete proofs with missing justifications.

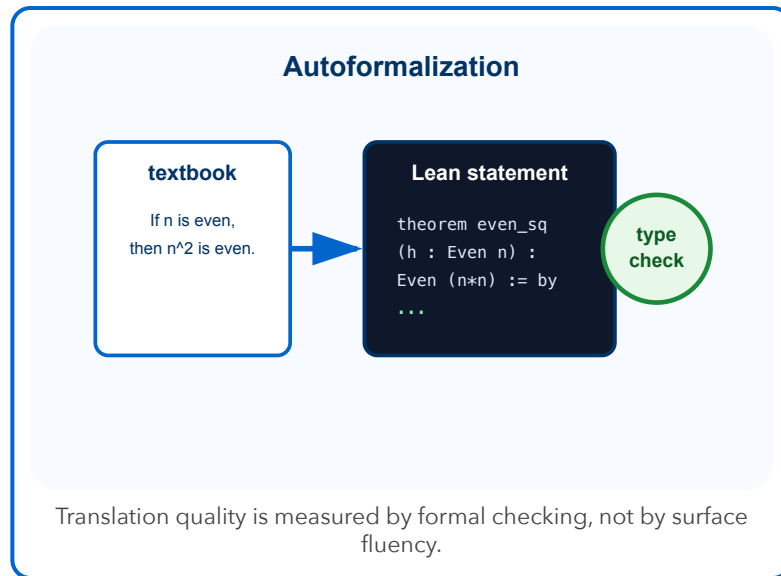
Metrics: precision/recall on planted errors, repair suggestion quality.



Project 4: Autoformalization

Question: How well can LLMs translate informal statements into Lean statements or proof skeletons?

- Start with elementary algebra, number theory, or real analysis statements.
- Record whether the generated Lean type-checks.
- Classify failures: wrong variables, missing assumptions, wrong library names.

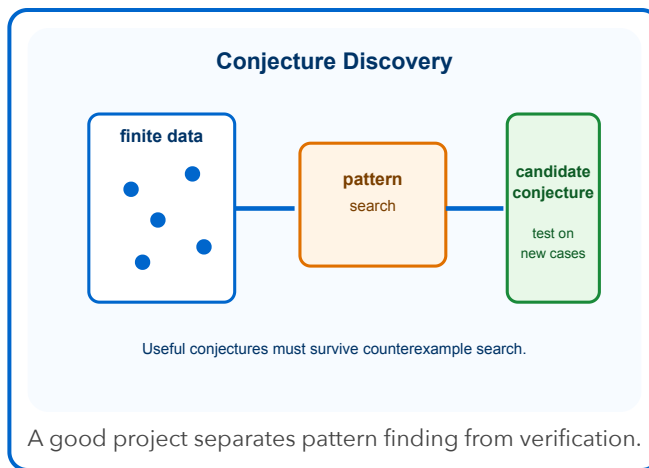


Project 5: Automated Conjecture Discovery

Question: Can simple ML or symbolic search propose plausible mathematical conjectures?

- Possible domains: integer sequences, graph invariants, finite groups, polynomial identities.
- Generate conjectures from small examples.
- Use held-out cases and counterexample search before trusting a pattern.

Metrics: novelty, truth rate on held-out examples, proofability, counterexample rate.

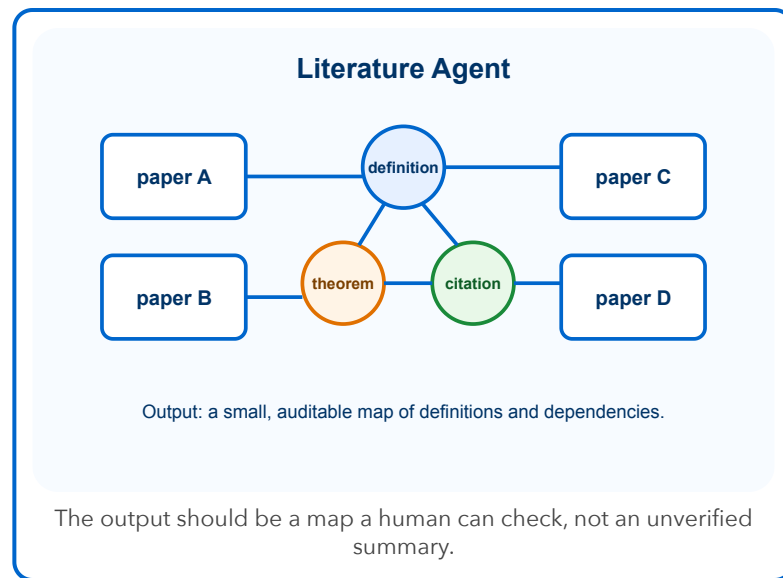


Project 6: Literature Agent

Question: Can an agent build a reliable map of a small mathematical topic?

- Choose a narrow topic: one theorem family or one technique.
- Extract definitions, assumptions, theorem dependencies, and citations.
- Audit every citation against the original source.

Metrics: citation accuracy, missing dependency rate, human correction time.

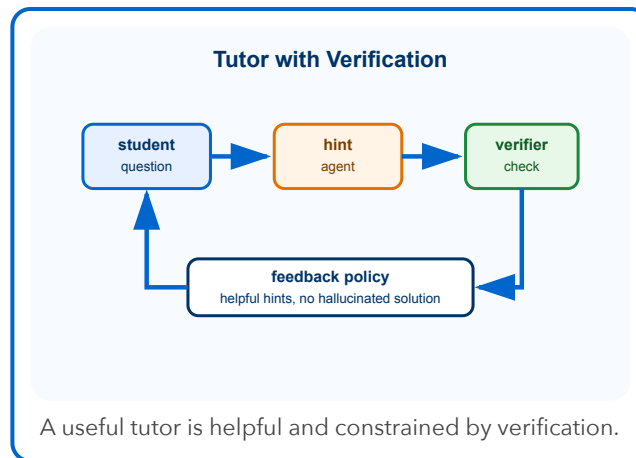


Project 7: Tutor with Verification

Question: Can a tutor agent teach problem solving while preventing hallucinated answers?

- **Hint first:** no direct full solution immediately.
- **Check claims:** use a checklist or Lean when possible.
- **Evaluate:** correctness and student usability.

Metrics: correct hints, avoided false claims, student success after hints.



Feasibility Test

- Can the first dataset be built quickly?
- Can the baseline be run without a large system?
- Can success be measured without relying only on opinion?
- Would a negative result still be meaningful?
- Can the scope be reduced if the toolchain becomes difficult?

Project Milestones

Time	Milestone
Launch	project seed chosen
Checkpoint 1	dataset/problem set and baseline ready
Checkpoint 2	first prototype or evaluation table
Checkpoint 3	failure analysis and revised method
Clinic	project clinic and presentation rehearsal
Presentations	final presentation and artifact review

What Counts as Success?

The project does not need to solve all of AI mathematics.

Clear experiment

The question is measurable.

Honest result

Positive or negative, the evidence is clear.

Useful artifact

Others can inspect prompts, data, logs, or code.

