College of AI – Preliminary Exam Syllabus

Rules

(1) **You have 3 hours.**

(2) **Total: 100 points.** In order to pass, PhD students should get at least **60** points.

(3) **Closed-book.** You may bring **five A4 pages** (hand-written or printed on both sides). **No** books, computers, or cellphones.

(4) You have to prepare for **one subarea**. There will be problems for each subarea in the exam. You only need to answer the problem in the subarea you choose.

General Requirements (85%)

1. Basic Algorithms

(1) **Sorting**

Bubble sort, insertion sort, merge sort, heapsort, quicksort

(2) **Basic Data Structures**

Binary search tree (CLRS Ch.12)

(3) **Basic Graph Problems**

a. BFS, DFS

b. Connected components

c. Directed acyclic graphs and topological order (CLRS 22.4, 22.5)

(4) **Minimum Spanning Tree**

a. Prim algorithm (KT 4.5, 4.6)

b. Can use MST to model and solve other problems

(5) **Dynamic Programming** (KT Ch.6, CLRS Ch.15) – can use DP to solve problems

(6) **Greedy Algorithms** (KT Ch.4, CLRS Ch.16)

a. Interval scheduling (KT 4.1)

b. Scheduling to minimize lateness (KT 4.2)

c. Huffman codes (KT 4.8)

(7) Basic Linear Programming (CLRS Ch.29) – know what a linear program is

2. Solving Problem by Search（AI Book Chapter 3）， Beyond classic search (AI Book Chapter 4)

3. Basic Machine Learning

(1) https://www.coursera.org/learn/machine-learning

(2) Things you need to know: Bias-Variance trade-off, overfitting， *k*-means, *k*-nn, PCA， Random Forest, AdaBoost, Gradient Boosting， Expectation Maximization (EM)，(Stochastic) Gradient Descent , SVM, kernel, lasso

4. Deep Learning

(1) https://cs231n.stanford.edu/2024/ (Lecture 1-13)

(2) **Deep Neural Network Architectures**: multi-layer perceptron, Autoencoders, Convolutional neural network, Transformers

(3) **Basic Components of Deep Learning** [see [DL book] Ch.7-15]: Backpropagation, Activation functions, Regularization for Deep Learning, Optimization Algorithms for Training Deep Model

(4) **Sequence Modelling**: Recurrent nets, Seq2seq model, Attention

(5) **Generative Models**: basic concepts of generative models; basic definition and idea of energy-based model, flow model, VAE, GAN, diffusion models

(6) **Science of Deep Learning (why does deep learning work)**: benefits of depth, properties of loss landscape (e.g., mode connectivity), training dynamics (e.g., grokking) [UDL book Ch.20]

5. Basic Reinforcement Learning

6. Basic Mathematics

Everything you learn from the calculus, linear algebra, and probability course, convex optimization [CO book].

Reference

[DL book] **Deep Learning.** Ian Goodfellow, Yoshua Bengio & Aaron Courville. https://www.deeplearningbook.org/

[AI Book] **Artificial Intelligence: A Modern Approach.** 3rd Ed. Stuart Russell & Peter Norvig

[RL book] **Reinforcement Learning: An Introduction** (2nd Ed.). Richard Sutton & Andrew Barto

[CV book] **Foundations of Computer Vision** – Antonio Torralba, Phillip Isola & William T. Freeman, MIT Press (2024)
[Math book] **Mathematics for Machine Learning.** A. Aldo Faisal, Cheng Soon Ong, Marc Peter Deisenroth
[ML book] **Pattern Recognition and Machine Learning.** Christopher Bishop. https://www.microsoft.com/en-us/research/wp-content/uploads/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf
[DL Course] **Stanford CS336 Language Modeling from Scratch.**
[UDL book] **Understanding Deep Learning.** Simons J.D. Prince.
[CO book] **Convex Optimization.** Stephen Boyd & Lieven Vandenberghe. https://web.stanford.edu/~boyd/cvxbook/
[KT book] **Algorithm Design.** Jon Kleinberg & Eva Tardos.

Sub-areas (choose one，15%)

Sub-area #1: AI Theory

1. **Applied Math**
    a. Spectral decomposition
    b. Eigenvalue decomposition
    c. Singular value decomposition
    d. Stochastic process and mixing
    e. Markov Chain Monte Carlo
2. **Machine Learning**
    a. PAC learning
    b. Bias-Variance tradeoff
    c. VC dimension
    d. SVM
    e. Decision Trees

Reference Books

[1] Blum, Avrim, John Hopcroft & Ravindran Kannan. Foundations of Data Science.

[2] Shalev-Shwartz, Shai & Shai Ben-David. Understanding Machine Learning: From Theory to Algorithms. Cambridge UP, 2014.

Sub-area #2: Computer Vision

1. **CV / Overview of Computer Vision**
   a. Image Sensing & Acquisition – CCD/CMOS sensors, optics and lens distortions, exposure control, noise models
   b. Camera Geometry & Calibration – pinhole model, intrinsic/extrinsic parameters, epipolar geometry, calibration workflows
   c. Core Vision Task Landscape – low-level (denoising, super-resolution), mid-level (feature extraction, segmentation), high-level (classification, detection, scene understanding)
   d. Benchmark Datasets & Metrics – MNIST, CIFAR, ImageNet, COCO, Cityscapes; evaluation metrics such as accuracy, IoU, mAP, F1

   [1] Computer Vision: Algorithms and Applications, 2 ed. – Richard Szeliski, Springer (2022)
   [2] Computer Vision Metrics: Survey, Taxonomy, and Analysis of Computer Vision, Visual Neuroscience, and Visual AI, 2 ed. – Scott Krig, Springer (2025)

2. **CV / Traditional Techniques**
   a. Traditional Filtering & Feature Extraction – Gaussian/median filters, Canny/Harris edges, SIFT & ORB descriptors
   b. Classical Recognition, Tracking & Scene Parsing – Bag-of-Visual-Words, HOG-SVM, Kalman/particle tracking
   c. Geometric & Motion Vision – optical flow, visual odometry

   [3] Feature Extraction and Image Processing for Computer Vision, 4 ed. – Mark S. Nixon & Alberto S. Aguado, Academic Press (2019)
   [4] Computer Vision: A Modern Approach, 2 ed. – David A. Forsyth & Jean Ponce, Pearson (2011)

3. **CV / Deep-Learning Approaches**
   a. CNN Architectures & Transfer Learning – modern convolutional backbones, feature reuse, fine-tuning across domains
   b. Detection, Segmentation & Dense Prediction – end-to-end object localization, pixel-wise labeling, panoptic vision tasks
   c. Vision Transformers & Spatiotemporal Models – attention-based encoders, hybrid CNN-Transformer designs, unified spatial-temporal reasoning
   d. Video Understanding & Temporal Modeling – integrated spatial-temporal feature extraction, multi-frame context, sequence-level recognition

   [5] Deep Learning – Ian Goodfellow, Yoshua Bengio & Aaron Courville, MIT Press (2016)

[6] Transformers for Natural Language Processing and Computer Vision, 3 ed. – Denis Rothman, Packt (2024)

4. **CV / Self-supervised & Unsupervised Learning**
    a. Contrastive Representation Learning – instance-level similarity objectives, multi-view feature alignment
    b. Masked & Predictive Image Modeling – context-driven reconstruction, masked region prediction, spatial forecasting
    c. Generative & Reconstruction-based Learning – latent variable modeling, image generation, and restoration signals
    d. Clustering, Pseudo-Labeling & Domain Adaptation – self-labeling cycles, category discovery, domain-shift adaptation

[7] A Cookbook of Self-Supervised Learning – Randall Balestriero et al., arXiv (2023)
[8] Self-Supervised Learning Crash Course: Unsupervised AI for NLP, Vision, and Beyond – Packt (2024)

5. **CV / Large Models & Multi-Modality**
    a. Vision–Language Foundation Models – large-scale joint pre-training on image–text pairs, unified embedding spaces, broad zero-/few-shot transfer across tasks
    b. Multimodal Generation – text-to-image and image-to-text synthesis, coherent cross-modal storytelling, iterative refinement workflows
    c. Modality Alignment – shared representations for retrieval, cross-modal matching and grounding, seamless information flow between visual, linguistic, and other sensor streams

[9] Multimodal Deep Learning – Cem Akkus et al., arXiv (2023)

6. **CV / 3-D Vision & Simulation**
    a. Depth & Geometry Estimation – stereo / monocular depth, multi-view reconstruction, surface fusion
    b. SLAM & Spatial Mapping – visual-inertial localization, loop closure, dense & sparse mapping
    c. Neural Rendering – neural radiance fields, differentiable rendering, novel-view synthesis
    d. Simulation – physics-based virtual worlds, domain randomization, sim-to-real data generation

[10] Computer Graphics: Principles and Practice, 3rd ed. — John F. Hughes et al., Addison-Wesley (2013)
[11] Fundamentals of Computer Graphics, 4th ed. — Steve Marschner & Peter Shirley (eds.), CRC Press (2016)

Sub-area #3: Natural Language Processing

1. **Basic Concepts**: Statistical NLP vs. linguistics; Syntax vs. semantics; Formal language theory; Statistical language modeling
2. **Formal Language Theory and Chomsky Hierarchy**: DFA & reduction; NFA; Regular expressions; Context-free grammar; Context-dependent grammar; Turing machines, P vs. NP
3. **Statistical Models**: Probabilistic grammars; Parsing; Earley Parser; Hidden Markov Models; Conditional Random Fields
4. **Linear Language Modeling and Learned Embeddings**: One-hot and N-gram features; Word2Vec; Arithmetic on learned embeddings
5. **Neural Language Models and Recurrent Networks**: Neural language models, teacher forcing; Likelihood, perplexity, entropy; Lossless compression; Vanilla RNN; Vanishing gradients and long-range dependencies; LSTM; GRU; RWKV; Mamba
6. **Attention and Transformers**: Neural attention to machine translation; BERT; T5 model; GPT model; Multi-token prediction; Belief States; Diffusion language models
7. **Scaling Laws and New Topics**: Instruction fine-tuning; RL from Human Feedback; Kaplan scaling law, chinchilla; Synthetic data, Phi models; Reasoning tokens, DeepSeek-R1; AI Agent systems; Long context modeling; Retrieval augmented generation

Reference Book

[1] Introduction to Automata Theory, Languages, and Computation – John Hopcroft
Jason Eisner NLP course: https://www.cs.jhu.edu/~jason/465/
[2] Foundations of Statistical NLP – Chris Manning

Sub-area #4: Reinforcement Learning

1. **Basic Concepts**: Markov Decision Process, value function, Bellman equation, exploration-exploitation trade-off; online vs. offline; model-free vs. model-based; planning vs. reactive policy; bandits vs. stateful RL; partial observability
2. **Offline-RL**: Behavior Cloning; Q-Learning; Deep Q-Learning
3. **Online-RL and Policy Gradient**: TRPO; PPO; GRPO; Actor-critic
4. **Model-Based Methods and Planning**: World Models; Dreamer; AlphaGo; MuZero
5. **Exploration**: Epsilon-greedy; Count-based; Elliptical Bonuses; Random Network Distillation; BYOL-Explore; MaxEnt RL; GFlowNet
6. **RL for LLMs**: RLHF; DPO; RL for reasoning

Reference Book
[1] [RL book] Reinforcement Learning: Theory and Algorithms. https://rltheorybook.github.io/
[2] [RL Lectures] Advanced Topics 2015, David Silver.
https://davidstarsilver.wordpress.com/teaching/
[3] [RL Book] Algorithms of Reinforcement Learning, Csaba Czepesvari.
https://sites.ualberta.ca/~szepesva/rlbook.html

Sub-area #5: Generative AI

1. **Theoretical Foundations**
    a. Transformer Architecture; Self-Attention; Mixture-of-Experts
    b. Diffusion Model
    c. Scaling Laws
2. **Large Language Models (LLMs)**
    a. Pre-training; Post-training; Alignment
    b. Long-Context Training
    c. Retrieval-Augment Generation
    d. Parameter-Efficient Fine-Tuning
    e. Safety, Red-Teaming, Guardrails
    f. Agent; Multi-Agent Systems
3. **Vision and Language**
    a. Contrastive Learning
    b. Multimodal Alignment
    c. Visual Instruction Tuning
    d. Masked/Generative Pretraining

Reference Book
[1] Vaswani, Ashish, et al. "Attention is all you need." NIPS 30 (2017).
[2] Kaplan, Jared, et al. "Scaling laws for neural language models." arXiv:2001.08361 (2020).
[3] Guo, Daya, et al. "Deepseek-r1: Incentivizing reasoning capability in LLMs via reinforcement learning." arXiv:2501.12948 (2025).
[4] Yang, An, et al. "Qwen3 technical report." arXiv:2505.09388 (2025).
[5] Liu, Haotian, et al. "Visual instruction tuning." NeurIPS 36 (2023): 34892-34916.
[6] Ho, J., Jain, A., & Abbeel, P. "Denoising diffusion probabilistic models." NeurIPS 33 (2020): 6840-6851.

[7] *https://huggingface.co/learn/diffusion-course/en/unit0/1*
[8] *https://huggingface.co/learn/computer-vision-course/en/unit4/multimodal-models/vlm-intro*

Sub-area #6: Physical AI

1. **Foundations & Modeling**
   a. Robot Perception & 3D Vision
   b. Multi-view geometry basics ([3DV Book] Ch.1-3)
   c. Point cloud processing ([Robotics Book] Thrun Ch.6, Barfoot Ch.8)
   d. Robotic Dynamics
   e. Rigid-body kinematics/dynamics ([Robotics Book] Siciliano Ch.2-3)
   f. Force/impedance control ([Robotics Book] Siciliano Ch.8-9)
2. **Algorithms & Planning**
   a. Motion planning & Trajectory planning ([Robotics Book] Siciliano Ch.4 & 12)
   b. Optimal control & MPC ([Robotics Book] Siciliano Ch.8-9)
   c. Reinforcement Learning & Imitation Learning
3. **Systems & Applications (30 %)**
   a. Autonomous Mobile Robots
   b. Robot Localization ([Robotics Book] Siegwart Ch.5)
   c. Navigation & obstacle avoidance ([Robotics Book] Siegwart Ch.6)
   d. Physical AI Toolchains
   e. ROS/Gazebo simulation
   f. NVIDIA Isaac Sim/Omniverse

Reference Book
[1] [3DV Book] Ma, Y. et al. An Invitation to 3-D Vision (2004).
[2] [Robotics Book] Siegwart, R. et al. Introduction to Autonomous Mobile Robots (2nd ed., 2011).
[3] [Robotics Book] Siciliano, B. et al. Robotics: Modelling, Planning and Control (3rd ed., 2010).
[4] [Robotics Book] Barfoot, T.D. State Estimation for Robotics (2022).
[5] [Robotics Book] Thrun, S. et al. Probabilistic Robotics (2005).

Sub-area #7: AI + Science

1. **Inverse Problem**
    a. Ill-posedness analysis
    b. Dimensionality reduction
    c. Regularization methods
    d. Prior modeling
2. **Physics-Informed Learning**
    a. Unfolding neural network
    b. Physics-informed neural networks
    c. Regularization flow
    d. Physics-constrained generative models
    e. Belief network
    f. Equivariant neural network
    g. Neural ordinary differential equation
3. **Scientific Computing**
    a. Scientific data and visualization
    b. Numerical analysis and simulation
    c. High performance computing
    d. Uncertainty quantification and optimal experimental design
    e. Probabilistic and causal inference
    f. Symbolic regression

Reference Book:

[1] Artificial Intelligence For Science: A Deep Learning Revolution. Alok Choudhary, Geoffrey C Fox, Tony Hey.

[2] Artificial Intelligence for Science (AI4S) Frontiers and Perspectives Based on Parallel Intelligence. Qinghai Miao and Fei-Yue Wang. https://link.springer.com/book/10.1007/978-3-031-67419-8

[3] Advanced Numerical Analysis: Machine Learning, Uncertainty Quantification, and Emerging Methods. Machine L.

[4] High-Performance Scientific Computing: Algorithms and Applications. Michael W. Berry, Kyle A. Gallivan, et al.

[5] Probabilistic Machine Learning: Advanced Topics (Adaptive Computation and Machine Learning series). Kevin P. Murphy.

[6] Causality: Models, Reasoning and Inference. Judea Pearl.

[7] Pattern Recognition and Machine Learning. Christopher M. Bishop. https://www.microsoft.com/en-us/research/wp-content/uploads/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf

---

Sub-area #8: AI System

1. **Architecture and Organization**

   1) Digital Logic

   a. Logic expressions and Boolean functions
   b. Can use De Morgan' s Law to simplify combination logic
   c. Sequential logic elements (flip-flops, latch)
   d. Basic timing constraints of sequential logic (synchronous circuits)
   e. Carry-ripple adder and array multiplier
   f. Signed and unsigned arithmetic
   g. Range, precision, and errors in floating-point arithmetic

   2) Computer Architecture

   a. Introduction to instruction set architecture, microarchitecture and system architecture
   b. Processor architecture – instruction types, register sets, addressing modes
   c. Memory hierarchy, latency and throughput
   d. Cache memories principles, replacement policies, and cache coherency
   e. Can use Amdahl' s law to prioritize design decisions
   f. GPU architecture and programming model
   g. NPU architecture and programming model

2. **Hardware-Software Codesign**

   1) Quantization

   a. Number representations for lightweight NN (BF16, FP8, INT4)
   b. Can calculate the scale and bias value in a quantized NN layer
   c. Concept of post-training quantization and quantization-aware training

   2) Pruning

   a. Understand structured and unstructured sparsity
   b. Sparse encoding formats (bitmask, run-length encoding, compressed sparse row)

c. Regularizations to construct sparsity in NN

3. **Networking and Interconnect**
    a. Background and history of networking and the Internet
    b. Network architectures
    c. Networks and protocols
    d. Distributed computing

Reference Book:

[1][Digital circuit book] CMOS VLSI Design: A Circuits and Systems Perspective (4th Edition), Neil H. E. Weste, David Money Harris

[2][Arch book] Computer Architecture: A Quantitative Approach (7th Edition), John L. Hennessy, David A. Patterson, Christos Kozyrakis

---

College of AI Graduate Program Committee – **June 2025**