

CSE 428 Computational Biology Capstone

Schedule and grading

1. Three presentations
2. course project presentation (60%), final report (40%)

Tentative Schedule

Date	Topic
4/1	Welcome/overview. Introduction to CSE428. (Sheng)
4/8	Project topic presentation (first half)
4/15	Project topic presentation (second half)
4/22	Working time
4/29	Mid-term project idea presentation (first half)
5/6	Mid-term project idea presentation (second half)
5/13	Working time
5/20	Working time
5/27	Final project presentation (first half)
6/3	Final project presentation (second half)

Goal in previous CSE428

- Work on a computational biology project
- Three steps
 - Pick a good paper
 - Reproduce that paper:
 - E.g., AI for protein structure prediction
 - Use the code repo released by the author
 - Propose new ideas to improve that paper (optional)

Goal for this year's CSE428

- Work on a computational biology project with significant help from AI/LLM
- Three steps
 - Pick a good paper
 - Discuss with LLM to understand why this is a good paper
 - Reproduce that paper
 - E.g., AI for protein structure prediction
 - Use the code repo released by the author
 - Use coding agent to help you reproduce the paper
 - Propose new ideas to improve that paper (optional)
 - Brainstorm with LLM regarding Feasibility and novelty


In my research group, we do 100% vibe coding

- Learn how to do vibe coding efficiently
- Understand the capability of different LLM and coding agent
- **This is impossible:** Prompt “please implement a new method that outperforms AlphaFold”
- This is possible:
 - Prompt “propose 10 ideas that might improve AlphaFold on metric X on dataset Y”
 - Prompt “implement these 10 ideas”
 - Prompt “examine whether one of these ideas outperform AlphaFold”
- This is a human-AI interaction process.
- Beyond vibe coding: use AI to write a report, plot figures, read papers, etc.

Today

- Logistics
- How to pick a good paper
- AI tools I suggested
- Specific projects I suggested

See office hours and other info on the course website.



The screenshot shows a web browser window with the address bar displaying `courses.cs.washington.edu/courses/cse428/`. The page header is a dark purple banner with the Washington University logo (a large 'W') on the left, the text 'PAUL G. ALLEN SCHOOL OF COMPUTER SCIENCE & ENGINEERING' in the center, and links for 'About Us' and 'Contact Us' on the right. The main content area is white. On the left side, there are two sections: 'Current Quarter' and 'Previous Quarters'. 'Current Quarter' lists 'Spring, 2026 (Wang)' and 'Prerequisites' as 'CSE 312', 'CSE 331', and 'CSE 332'. 'Previous Quarters' lists a list of past offerings from Spring 2025 down to Spring 2009, with instructors in parentheses. The right side of the page features the course title 'CSE428: Computational Biology Capstone' in a large, bold font. Below the title is the 'Catalog Description', 'Prerequisites', 'Credits', and 'ABET Outcomes'.

CSE428: Computational Biology Capstone

Catalog Description: Designs and implements a software tool or software analysis for an important problem in computational molecular biology. Prerequisite: CSE 312; CSE 331; CSE 332.

Prerequisites: CSE 312; CSE 331; CSE 332.

Credits: 5.0

ABET Outcomes:

This course contributes to the following [ABET](#) outcomes:

(1) an ability to identify, formulate, and solve complex engineering problems by applying principles of engineering, science, and mathematics

(2) an ability to apply engineering design to produce solutions that meet specified needs with consideration of public health, safety, and welfare, as well as global, cultural, social, , and economic factors

(3) an ability to communicate effectively with a range of audiences

(4) an ability to recognize ethical and professional responsibilities in engineering situations and make informed judgments, which must consider the impact of engineering solutions in global, economic,

- course project presentation (60%), final report (40%).

Very tight schedule this year! Only 9 lectures.


Date	Topic
3/30	Welcome/overview. Introduction to CSE428. (Sheng)
4/6	Project topic presentation (first half)
4/13	Project topic presentation (second half)
4/20	Working time
4/27	Mid-term project idea presentation (first half)
5/4	Mid-term project idea presentation (second half)
5/11	Working time
5/18	Final project presentation (first half)
6/1	Final project presentation (second half)

How to pick a good paper

- This is the part that AI can offer the least help because this depends on your interests
- Let's first see different areas of computational biology

Three categories of biomedical research

Mechanisms of biology



Drug discovery



Daily diagnosis



Generate scientific hypothesis Generate molecular structure Model multi-modal patient data

AI Scientist

Generation

Multi-modal understanding





Basic science


Real-world applications

Human health data are inherently multi-modal

Complication of diabetes




Machine learning problem setting




1D data: 3 billion long DNA sequence

Pathology Dermoscopy




Funds


X-ray



2D data: different images




3D data: videos and 3D images



4D data: time-series of multiple visits


Patient embedding



Patient notes

Specialty Consultation
Subjective: Patient presents with a rash on their arms and legs that has been present for the past week. They report itching and discomfort, but no pain. Patient states they have not changed any skincare products recently.
Objective: Multiple raised red patches were noted on both arms and legs. Patches appear to be slightly swollen and warm to the touch.
Assessment: Allergic reaction or contact dermatitis.
Plan: Recommend avoiding any potential irritant/allergens that may have caused the rash. Topical corticosteroid cream to alleviate itching and inflammation was given. A follow-up appointment in two weeks to monitor improvement.


Data structure for each scale: protein



A sequence of amino acids/nucleic acids -> A sequence of word/character
NLP methods (edit distance, LSTM, BERT)


Computational challenge: modeling the order in the sequence

Data structure for each scale: network



Computational challenge: interaction, synergistic effect

Data structure for each scale: cell



High-dimensional, noisy, large-scale

Data structure for each scale: tissue





Image analysis, lack of high-quality annotations



Data structure for each scale: organism





Protein
(1 nm)



Protein complexes (function)
(10-100nm)



Cell
(1-10 μ m)



Tissue
(100 μ m to 100 mm)



Complex organism
(> 1cm)



Disease mechanisms -> Multimodality
Integration of information from sequences, networks,
images and matrixes

Multi-modality and heterogeneous

How to pick a good paper

- This is the most important thing
 - One researcher spends 3 hours per day to read paper
- Read new papers (2023-) and papers published in top journals/conferences
 - Recomb, ISMB
 - Nature Methods, Nature Biotechnology, Nature Medicine, Nature Genetics
 - Nature Biomedical Engineering, Nature Cancer
 - Nature Machine Intelligence, Nature Communications, Nature Computational Science
 - Nature, Science, Cell
- Tips for how to read papers (as a beginner)
 - Focus less on the background, introduction and motivation
 - Focus more on the problem setting (input, output)
 - Focus more on the datasets for evaluation (size, accessible, simulated data or real data)
- The more you read, the faster you will read.
- Ultimate goal: you can “predict” what method/dataset/evaluation this paper will do after reading the abstract

Pick the problem to work on

- Don't be ambitious!
- Only work on problem that an existing paper has worked on and you can fully understand that paper
- Only work on problem that you can clearly evaluate

Pick the problem to work on

- Pick a paper from Profs from Allen School
 - Su-In Lee, Sara Mostafavi, Sheng Wang, Bill Noble
- Read at least three papers in the first week
 - Each student will do a presentation next week
 - Significance
 - Innovation
 - Methods and evaluation
 - Why do you like this paper?

Borrow ideas from other domains

- From less novel to very novel
 - Use a method in one comp bio problem (e.g., protein structure) to solve another comp bio problem (e.g., chromatin structure)
 - Use a new ML technique from CV/NLP to solve a comp bio problem (CSE428 will focus on this)
 - Use a method in other subjects (quantum computing, theoretical physics) to solve a comp bio problem

How to borrow ideas from other domains

- Find the commonality by understanding the problem setting
 - Problem A in computer vision has the same problem setting as Problem B in comp bio.
 - Talk to people work in CV/NLP

Experiments

1. Don't implement your ideas first
2. First, reproduce the results from baselines
3. Test baselines on new datasets and examine the performance
4. If the performance is bad, figuring out the reason

Uniqueness of comp bio research

- Interdisciplinary subject
 - Communication is very important
 - Understand biologists/doctors' needs
 - Clearly explain our methods
- New subject
 - Lots of opportunities
 - Less well-defined benchmarks

AI tools I use


- Cursor + Claude code
- Let's do a quick demo.

How we finish a Nature Methods paper in four weeks

1. Prof. Noble has a proposal for a new idea (AI cannot do it)
2. Sheng rewrites this proposal a format that AI can understand
 - Mostly simplify instead of adding more details!
 - AI needs flexibility
3. Use GPT to propose detailed plan based on the proposal
4. Use Cursor and Claude to implement it.
5. Ask Claude code to summarize experiment results
6. Use Gemini (Nano Banana) to plot figures
7. Prof. Noble writes the paper.



Prof. Bill Noble
UW Genome Sciences



Tangqi Fang, PhD student
UW CSE

Proposal we sent to AI model


```
1 # Run Foundation Model Proposal
2 ## Objective
3 Learn a run-level representation from raw mass spectrometry data and evaluate whether it captures a latent experimental state beyond pooled spect
4
5 ## Hypothesis
6 A run-level embedding encodes global experimental structure that cannot be recovered by aggregating independent spectrum-level embeddings.
7
8 ## Data
9 - Raw runs:
10   ~/datadrive/local_data/run/PRIDE_processed
11 - Use only:
12   - Raw spectra (MS1 + MS2)
13   - Scan-level metadata (e.g., retention time)
14   - File path / project structure
15
16 ## Baselines
17 1. Spectrum-level embedding + pooling
18 2. Run-level statistical features
19
20 ## Data Processing
21 - Separate MS1 and MS2 spectra.
22 - For each run and level (MS1/MS2), perform RT-stratified sampling: split the RT range into CAP bins and select the spectrum with the highest TIC
23 - Represent each spectrum via log-m/z binning into N_MZ_BINS between MIN_MZ and MAX_MZ using nearest-bin accumulation (fixed-length vector).
24 - Intensity preprocessing is strictly: sqrt → log(1+x); do not apply per-spectrum max or min-max normalization.
25
26 ## Model
27 - Stage 1: Pretrain spectrum encoder using Transformer + MAE
28 - Stage 2: Aggregate spectrum embeddings into run embedding using Transformer and retention time as positional embedding
29
30 ## Evaluation Tasks
31 1. **Absolute RT Prediction**
32   Predict retention time without per-run calibration.
33 2. **Tissue Label Prediction**
34   Supervised classification using run embeddings.
35 3. **Project Classification / Clustering**
36   Evaluate alignment with PRIDE project structure.
```

A multi-agent system (designed by AI)


1. Data downloading
 1. Input: proposal, url of the data
 2. Output: raw data saved on the server
 3. Very hard for AI. We did it manually
2. Data processing
 1. Input: raw data
 2. Output: preprocess data, split into train, test.
 3. A little bit hard for AI. But we need to manually examine.
3. AI modeling
 1. Input: processed data, successful metrics
 2. Output: trained AI model
 3. Very easy for AI. We don't need to examine anything. AI will design architecture and algorithms
4. Downstream evaluation
 1. Similar to 1-3. We decompose the problem into data downloading, data processing and AI modeling.

Four paradigms in AI for Medicine


Deep learning
One model for one task




Foundation model
One model for all tasks



Multi-modal foundation model
One model takes different inputs





Multi-agent
Integrate multiple foundation models



Medicine is inherently multi-modal

Complication of diabetes




Today's talk: 3 parts


1. Pathology foundation model
2. 3D retinal foundation model
3. A multi-modal foundation model integrating 9 imaging modalities




Multi-modal foundation model (2023)
One model takes different inputs




Foundation model (2022)
One model for all tasks



Multi-agent (2024)
Integrate multiple foundation models





Pathology images are too large for existing AI models



Pathology images
100k by 100k pixels

Pathology images are too large for existing AI models





Pathology images
100k by 100k pixels

As large as a tennis court



Pathology images are too large for existing AI models



Pathology images
100k by 100k pixels

As large as a tennis court




As large as a tennis ball





Images handled by existing AI models
256 by 256 pixels

Long-context modeling: ChatGPT is good at understanding long documents




Long-context modeling: ChatGPT is good at understanding long documents



Use LongNet to model a **long sentence** of small images from a pathology slide

Tile-level encoder to capture local patterns

Slide-level encoder to capture the pattern in the whole slide



A self-supervised learning framework based on DinoV2 and LongNet

GigaPath: A whole-slide foundation model for pathology

nature

Explore content ▾ About the journal ▾ Publish with us ▾

[nature](#) > [articles](#) > [article](#)

Article | [Open access](#) | Published: 22 May 2024


A whole-slide foundation model for digital pathology from real-world data

Hanwen Xu, Naoto Usuyama, Jaspreet Bagga, Sheng Zhang, Rajesh Rao, Tristan Naumann, Cliff Wong, Zelalem Gero, Javier González, Yu Gu, Yanbo Xu, Mu Wei, Wenhui Wang, Shuming Ma, Furu Wei, Jianwei Yang, Chunyuan Li, Jianfeng Gao, Jaylen Rosemon, Tucker Bower, Scohee Lee, Roshanthi Weerasinghe, Bill J. Wright, Ari Robicsek, Brian Piening, Carlo Bifulco , Sheng Wang  & Hoifung Poon  — [Show fewer authors](#)

[Nature](#) **630**, 181–188 (2024) | [Cite this article](#)

200k model downloads every month
(Hugging Face)!

Media coverage: [Forbes](#), [Yahoo](#), [Becker's hospital review](#), [Fierce biotech](#), [CTOL digital solutions](#), [HIT consultant](#), [GeekWire](#), [Cosmic log](#), [HealthXL](#), [RamaOnHealthcare](#), [Providence](#), [nikkei](#), [cryptorank](#), [deeptech](#)




Hanwen Xu
U of Washington



Naoto Usuyama
Microsoft Research



Carlo Bifulco
Providence




Sheng Wang
U of Washington



Hoifung Poon
Microsoft Research


One (foundation) model fits all

30,000 patients from 28 hospitals
140k slides, 15 cancer types




GigaPa

Best performance on 25 out of 26 tasks



Cancer subtype classification using pathology images

Ovarian



?

EOV

CCO





HGSO

LGSCO


MO

OC

■ GigaPath ■ HIPT ■ CtransPath ■ REMEDIS




Biomarker prediction using pathology images: best performance on 17 out of 18 biomarker predictions



KRAS?


Mutate

Non-



Directly predicting treatment is too difficult

Different cancer treatment




Patient's pathology image

?

Clinical trial




Radiotherapy




Chemotherapy



Targeted therapy



Immunotherapy



Bone marrow transplant



Chain-of-Thought: Decompose a complicated task into many subtasks

Standard GenAI

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. ❌

Chain-of-Thought GenAI

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?


A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅

Our idea: use Clinical Guideline as the Chain-of-Thought




Our idea: use Clinical Guideline as the Chain-of-Thought



National
Comprehensive
Cancer
Network®

NCCN Guidelines Version 7.2024 Non-Small Cell Lung Cancer




Human-AI collaboration: Experts derive the guideline. AI makes decision on each branch.
Future implication: AI model as a clinical lab test

Today's talk: 3 parts


1. Pathology foundation model
2. 3D retinal foundation model
3. A multi-modal foundation model integrating 9 imaging modalities




Foundation model (2022)
One model for all tasks



Multi-modal foundation model (2023)
One model takes different inputs




Multi-agent (2024)
Integrate multiple foundation models



Lots of medical images are 3D




General-domain images: 2D




3D biomedical images: CT, MRI, OCT

Disease pattern is very small and hard to detect

Age-related macular degeneration



A GenAI framework for 3D medical imaging: MAE + FlashAttention





Convert a 3D volume to a **long sentence** of small cubes

The first 3D OCT foundation model trained from 26,605 patients

Best performance on all 29 tasks


— OCTCube — RETFound (all) — RETFound (center)




Work done by Zixuan Liu
U of Washington



Dr. Miao Zhang
Genentech




Dr. Cecilia Lee
UW ophthalmology




Dr. Aaron Lee
UW ophthalmology


3D model offers accurate and interpretable predictions for retinal diseases




3D model offers accurate and interpretable predictions for retinal diseases



OCT slice




3D model offers accurate and interpretable predictions for retinal diseases




OCT slice

RETFound




3D model offers accurate and interpretable predictions for retinal diseases




OCT slice

RETFound

Our method




3D model is better at generalization: Cross-cohort prediction at Duke, University of Minnesota, Johns Hopkins, NYU



Training




Test




3D model enables cross-device prediction

Heidelberg Spectralis




Topcon Masestro2



3D model enables cross-device prediction

■ OCTCube ■ RETFound (all) ■ RETFound (center)

Heidelberg Spectralis



Pretrai




Heidelberg Spectralis


Fine-



Topcon Masestro2



Topcon Masestro2



Pretrai




Heidelberg Spectralis




Fine-



Zeiss cirrus




Predict systemic diseases





One-year early prediction of diabetes

1st visit: **2D model** cannot identify diabetes




1st visit: **3D model** detects diabetes signal





One-year early prediction of diabetes


1st visit: **2D model** cannot identify diabetes







1st visit: **3D model** detects diabetes signal



1 year later



2D model detects diabetes




Today's talk: 3 parts


1. Pathology foundation model
2. 3D retinal foundation model
3. A multi-modal foundation model integrating 9 imaging modalities




Foundation model (2022)
One model for all tasks




Multi-modal foundation model (2023)
One model takes different inputs




Multi-agent (2024)
Integrate multiple foundation models




Multi-modality analysis in GigaPath and OCTCube



Medical report + Pathology image



OCT image + IR image



Existing cross-modal framework is
Limited to two modalities






Image from CLIP-ReIdent

Our solution for integrating all image modalities


- Build one foundation model for each modality, then project different models into the same space
- What should be the Esperanto for medicine?




BiomedParse: use human language as the Esperanto to integrate medical imaging modalities




Liver in abdomen CT






Kidney in abdomen CT



Pancreas in abdomen CT



Liver vessel in abdomen CT Inflammatory cells in pathology Neoplastic cells in pathology



A foundation model for joint segmentation, detection and recognition across 9 modalities

nature methods

[Explore content](#) ▾ [About the journal](#) ▾ [Publish with us](#) ▾ [Subscribe](#)

[nature](#) > [nature methods](#) > [articles](#) > [article](#)

Article | Published: 18 November 2024

A foundation model for joint segmentation, detection and recognition of biomedical objects across nine modalities

[Theodore Zhao](#), [Yu Gu](#), [Jianwei Yang](#), [Naoto Usuyama](#), [Ho Hin Lee](#), [Sid Kiblawi](#), [Tristan Naumann](#),
[Jianfeng Gao](#), [Angela Crabtree](#), [Jacob Abel](#), [Christine Mounq-Wen](#), [Brian Piening](#), [Carlo Bifulco](#), [Mu Wei](#)
✉, [Hoifung Poon](#) ✉ & [Sheng Wang](#) ✉


[Nature Methods](#) (2024) | [Cite this article](#)



Theodore Zhao



Aiden Gu



Sheng Wang
U of Washington




Mu Wei



Hoifung Poon

Available as open-source model on Azure AI and Hugging Face

Demo I: talk to the AI model to find lung nodule




Demo I: talk to the AI model to find lung nodule




Demo2: talk to the model to find all cell types

Input Image





Parsing Results




Text prompt

g

Output text

Go!

Demo2: talk to the model to find all cell types




The interface displays a histological image on the left under the label "Input Image". A large play button is centered between the input and output images. On the right, under the label "Parsing Results", the same image is shown with colored overlays and labels: "connective tissue cells" (blue), "inflammatory cells" (orange), and "neoplastic cells" (green).

Below the images, there is a "Text prompt" section with the input text: "get all cell types". To the right, the "Output text" section displays the model's response: "I found the following objects in the image: connective tissue cells, inflammatory cells, neoplastic cells." A blue "Go!" button is located at the bottom center.


0.2 second to segment and label all cells with >
90% consistency

Pathologist annotation



 Tumor region  Lymphocyte  Stroma




BiomedParse prediction



 Neoplastic cells  Inflammatory cells
 Connective tissue cells


Data-raising from Collaborators at UW School of Medicine

27k OCT images for retinal diseases




Drs. Cecilia and Aaron Lee (Ophthalmology)

50k Mammogram for breast cancer




Dr. Christopher Lee (Radiology)

100k CT for heart transplant



Dr. Shin Lin (Cardiology)

190k EM image for kidney




Dr. Behzad Najafian (Pathology)




Dr. Mehmet Kurt (ME)




100k brain MRI for stroke, Parkinson, brain cancer




Dr. Ricky Wang (BioE)




5k FA for glaucoma



Drs. Donald Chi, Amy Kim (Dentistry)



50k dental panoramic X-ray



Dr. Jay Pal (Surgery)



300k echo for heart failure




Drs. Nathan Cross and Paul Kinahan (Radiology)




120k spinal MRI for back pain

Four paradigms in AI for Medicine


Deep learning
One model for one task




Foundation model
One model for all tasks




Multi-modal foundation model
One model takes different inputs




Multi-agent
Integrate multiple foundation models



A Microsoft 365 for cancer diagnosis




Molecular tumor board: From multi-modality to multi-agent




Four paradigms in AI for Medicine


Deep learning
One model for one task




Foundation model
One model for all tasks



Multi-modal foundation model
One model takes different inputs



Multi-agent
Integrate multiple foundation models




Three lessons we learnt from GenAI for Medicine

- ✓ Medical foundation models are accurate

Three lessons we learnt from GenAI for Medicine


- ✓ Medical foundation models are accurate
- ✓ Medical foundation models are heterogenous

Data sources of three existing pathology models




Three lessons we learnt from GenAI for Medicine

- ✓ Medical foundation models are accurate
- ✓ Medical foundation models are heterogenous




Data sources of three existing pathology models



- ✓ **Less than 0.001% existing data** is used by even the largest medical foundation model
 - CT: 80m per year in US **vs.** 100k (10 years) by largest CT model
 - Pathology: 100m slides per year in US **vs.** 170k slides (10 years) by largest pathology model

Fifth paradigm: a world model


build one GenAI model using medical data all over the world



Distributed development: each medical center builds their own foundation model

Fifth paradigm: a world model

build one GenAI model using medical data all over the world



Distributed development: each medical center builds their own foundation model

Mixture-of-experts: a lightweight model integrates all these models

Fifth paradigm: a world model


build one GenAI model using medical data all over the world

What we need to build a world model?

- Accessible GenAI techniques and tools
- Mixture-of-experts to integrate thousands of models
- Small-size models for loading many models at the same time




How will the world model help?

- Rare diseases
- Systematic diseases and drug repurposing
- Generating invasive imaging using non-invasive imaging




Data-raising from Collaborators at UW School of Medicine

27k OCT images for retinal diseases




Drs. Cecilia and Aaron Lee (Ophthalmology)

50k Mammogram for breast cancer




Dr. Christopher Lee (Radiology)

100k CT for heart transplant



Dr. Shin Lin (Cardiology)

190k EM image for kidney




Dr. Behzad Najafian (Pathology)




Dr. Mehmet Kurt (ME)




100k brain MRI for stroke, Parkinson, brain cancer




Dr. Ricky Wang (BioE)



5k FA for glaucoma



Drs. Donald Chi, Amy Kim (Dentistry)



50k dental panoramic X-ray



Dr. Jay Pal (Surgery)



300k echo for heart failure




Drs. Nathan Cross and Paul Kinahan (Radiology)





120k spinal MRI for back pain

A Multimodal Biomedical Foundation Model Trained from Fifteen Million Image–Text Pairs.






A Multimodal Biomedical Foundation Model Trained from Fifteen Million Image–Text Pairs.



Question:	Are there multiple or just 1 metastatic focus?	What are the hyperdensities on the periphery of the image?	What is the biological sex of the patient?
Answer:	one	ribs	female
MEVF:	right chest ✗	storage of urine ✗	inflammation ... ✗
QCR:	no ✗	intestine ✗	treat brain diseases ... ✗
PubMedCLIP:	yes ✗	spinal cord ✗	nodule ✗
BiomedCLIP:	right lobe of liver ✗	ribs ✓	female ✓

A Multimodal Biomedical Foundation Model Trained from Fifteen Million Image–Text Pairs.



Question:	Are there multiple or just 1 metastatic focus?	What are the hyperdensities on the periphery of the image?	What is the biological sex of the patient?
Answer:	one	ribs	female
MEVF:	right chest ✗	storage of urine ✗	inflammation ... ✗
QCR:	no ✗	intestine ✗	treat brain diseases ... ✗
PubMedCLIP:	yes ✗	spinal cord ✗	nodule ✗
BiomedCLIP:	right lobe of liver ✗	ribs ✓	female ✓


Poisoning medical knowledge using large language models

Received: 11 October 2023

Accepted: 15 August 2024

Published online: 20 September 2024

Junwei Yang¹, Hanwen Xu², Srubhi Mirzoyan¹, Tong Chen², Zixuan Liu²,
Zequn Liu², Wei Ju¹, Luchen Liu², Zhiping Xiao²✉, Ming Zhang¹✉ &
Sheng Wang²✉



nature machine intelligence

Article

<https://doi.org/10.1038/s42256-024-00876-w>

A bioactivity foundation model using pairwise meta-learning

Received: 2 November 2023

Accepted: 3 July 2024

Published online: 14 August 2024

Bin Feng^{1,2}, Zequn Liu², Nanlan Huang¹ , Zhiping Xiao³ ,
Haomiao Zhang¹, Srбуhi Mirzoyan², Hanwen Xu³, Jiaran Hao¹,
Yinghui Xu^{1,4} , Ming Zhang²  & Sheng Wang³ 


A bioactivity foundation model using pairwise meta-learning

Received: 2 November 2023

Accepted: 3 July 2024

Published online: 14 August 2024

Bin Feng^{1,2}, Zequn Liu², Nanlan Huang¹, Zhiping Xiao³✉, Haomiao Zhang¹, Srбуhi Mirzoyan², Hanwen Xu³, Jiaran Hao¹, Yinghui Xu^{1,4}✉, Ming Zhang²✉ & Sheng Wang³✉

a

1 A generalizable Hi-C foundation model for chromatin architecture,
2 single-cell and multi-omics analysis across species

3 Xiao Wang^{*1,2}, Yuanyuan Zhang^{*3}, Suhita Ray⁴, Anupama Jha¹, Tangqi Fang², Shengqi
4 Hang², Sergei Doulatov^{†4}, William Stafford Noble^{†1,2}, and Sheng Wang^{†2}


5 ¹Department of Genome Sciences, University of Washington, Seattle, WA, USA

6 ²Paul G. Allen School of Computer Science and Engineering, University of Washington,
7 Seattle, WA, 98105, USA

8 ³Department of Computer Science, Purdue University, West Lafayette, IN, 47907, USA

9 ⁴Division of Hematology and Oncology, University of Washington, Seattle, WA, 98105, USA

C Pre-training of HiCFoundation



d HiCFoundation applications

