

Vision and Language

AthensNLP Summer School 2026




Desmond Elliott
University of Copenhagen




Working Definition

Multimodal models jointly processes information from two or more input modalities, e.g. images and text, speech and video, etc.




Why Multimodality?

- Humans ground conceptual knowledge in modality processing systems in the brain
- Evidence that grounding activates similar brain regions for different input modalities




Multimodality reduces ambiguity



You Cannot Learn Language From

- The radio without grounding
(lack perception)
- The television without actions
(lack embodiment)
- Without interacting with others
(lack social)





The Five Major Areas

- **Representation:** how to convert raw inputs into a usable format
- **Translation:** transform from one modality to another
- **Alignment:** predict relationships between elements across modalities
- **Fusion:** join features from modalities to support prediction
- **Co-learning:** transferring knowledge from one modality to another

Representation

- Great deal of work over the last decade, from HOG features in the early 2000s to features from models like CLIP or DINO in the 2020s.





Dalal & Triggs. CVPR 2005. Histograms of oriented gradients for human detection.

Radford et al. ICML 2021. Learning transferable visual models from natural language supervision.

Carom et al. ICCV 2021. Emerging properties in self-supervised vision transformers.



Translation

- Explosion of end-to-end neural network models since the mid 2010s




Alignment

- Important for self-supervised learning, and also for cross-modal retrieval or phrase grounding





Fusion

- Early work studied the differences between early and late fusion.
- Multi-head self-attention now provides model-based fusion.



Co-learning

- Zero-shot transfer across modalities, or using visual grounding to improve language models on text-only tasks.



Socher et al. (2013). Zero-shot learning through cross-modal transfer. NeurIPS.

Tan & Bansal. (2020). Vokenization: Improving Language Understanding with Contextualized, Visual-Grounded Supervision. EMNLP.


Roadmap

1. Datasets and Tasks for Multimodal Learning

 FoodieQA: A Multimodal Dataset for Fine-Grained Understanding of Chinese Food Culture

 Kaleidoscope: In-language Exams for Massively Multilingual Vision Evaluation

2. Data Representation



3. Modelling Techniques

 Efficient Test-Time Scaling for Small Vision-Language Models

4. Understanding Multimodal Models

 LatentLens: Revealing Highly Interpretable Visual Tokens in LLMs

5. Future Directions

1. Datasets and Tasks for Multimodal Learning

There Are No New Ideas in AI... Only New Datasets

LLMs were invented in four major developments... all of which were datasets



JACK MORRIS


APR 09, 2025

Two Basic Types of Resource

1. Mined from existing data sources



2. Created from scratch



Four Dataset Creator Stereotypes

Data Miner



Large-scale
extraction

Detectorist



Focused search
for artefacts

Permaculturist



Minimal
intervention

Baroque



Painstaking
design

Five Key Factors

1. Scope

What type of data are you aiming to collect?

2. Annotator Relationship

How are you working together?

3. Images

What type of images?

4. Texts

What type of texts?

5. Binding

How tightly related is the multimodal data?



Use Cases



Multi30K (2015)



1. The two men on the scaffolding are helping to build a red brick wall.
2. Zwei Maurer mauern ein Haus zusammen.



FoodieQA (2024)

以下菜品是哪个地区的特色菜？
Which **region** is this food a specialty?



- ☒ A 江苏 (Jiangsu)
☐ B 京津 (Beijing & Tianjin)
☐ C 香港 (Hong Kong)
☐ D 广西 (Guangxi)


MaRVL (2021)



(b) Görsellerden birinde dizlerinde kanun bulunan birden çok insan var. ("In one of the images, there are multiple people with qanuns on their knees.", concept: *Kanun (çalgı)* (QANUN, a popular instrument in Turkey), label: TRUE)



Kaleidoscope (2025)



University Entrance Exam

Physics

Question:

گلوله‌ای مسیری مطابق شکل را طی می‌کند. کار نیروی وزن از A تا B تا C چند برابر کار نیروی وزن از B تا C است؟


Options:

- A. 1.5
B. 1.25
C. 1.2
D. 2



1. Scope

- What type of data are you aiming to collect?
- Does your dataset come from a broad collection of concepts / domains, or is it a deep dive into a specific subject matter?



Scope: Task Types

- Sequence generation: $P(x|v)$ or $P(y|v)$
 - Captioning, zero-shot image classification, MCQA, image generation
- Classification: $P(y|x, v)$
 - VQA, Visually-grounded Reasoning
- Ranking and Alignment: **Distance**(x, v)
 - Image↔Text Retrieval
Referring Expression Localization


Multi30K: Replicate



A brown dog is running
after the black dog.

Ein schwarzer und ein
brauner Hund rennen auf
steinigem Boden
aufeinander zu

MaRVL: Concepts beyond ImageNet



FoodieQA: Fine-grained Chinese Food




Kaleidoscope: Exams

Geography

Question #5357

'A' চিত্রের প্রদর্শিত ঘটনাটি কী?



1) সঞ্চার

2) অগ্নিবৃপাত

3) ক্ষয়


4) উত্তরণ

Social Sciences | Bengali | Bangladesh

Mathematics

Question #9326

Alice a beaucoup de blocs, tous de dimensions 1 cm x 2 cm x 4 cm. Elle en met le plus possible dans une boîte de dimensions 4 cm x 4 cm x 4 cm sans qu'aucun ne dépasse. Combien a-t-elle rangé de blocs dans la boîte ?



1) 6

2) 7


3) 8

4) 9


STEM | French | France

2. Annotator Relationship

- What is the relationship between the collectors and the team?
- Are you working together towards a shared goal, or are they hired?



Multi30K: Hired Workers



Source: machine translated extract from crowdworker discussion forum about our crowdsourcing task.

Multi30K: Discussions on Tasks and Task Initiators

D#

2016 07:13

So for me they are very popular as we get here in Austria hardly good task or surveys that it is only an advantage for us that there is the image description 😅

I do after every 10th task a smoke break otherwise also it gets too steep to me

And today also again managed 50 of them, which is a great month 😊

D#

2016 09:41

So, I got the today 40 'managed only 8 minutes rest between. Unfortunately, I have come to the images only when steel Blue had already announced 46 minutes earlier, so time was short.

I now have a "system" designed not to be mad. 25 Task / 15 min pause / 25 task.

Should really go all 50, so you IF time begins, of course. ^^

D#;

2016 15:24

after I have just considered when cooking, as I would probably describe this cutlets in the pan, I'll let it go with the pictures describe today prefer 😊

Multi30K: Speculations

📄#:

2016 07:14

“



I think it's about how different people perceive images ... so what is the thing that strikes nearly all or where the focus is. Possibly even something for an AI ... laughing Since many photos which you describe I have for example, the woman at the end of the stairs lying etc.

Yes good chance when I look so what he has been previously employed. I find something fascinating 😊

way, I have just the absolutely is a dark image of a man apparently in a halfpipe and jumps, but you can not identify absolutely with what if skates or board or something else ... I hate something xD


Kaleidoscope: Collaboration

- Open-science collaboration with an incredible community of early-career scholars
- Co-authorship offered in exchange for collecting data above a threshold




3. Images


- What type of images are in your resource?
- From which sources are you collecting them?
- What are the licenses of the images?
- How are you protecting PII?



FoodieQA: Private Images




Kaleidoscope: Eight Image Types




4. Texts

- What type of text are you collecting?
- Where are you getting the text from?
- Which languages are you covering and why?




Kaleidoscope: Natural Questions

- The texts were created by exam-board professionals for educational purposes



Consider the Deterministic Finite-state Automaton (DFA) A shown below. The DFA runs on the alphabet $\{0, 1\}$, and has the set of states $\{s, p, q, r\}$, with s being the start state and p being the only final state. Which one of the following regular expressions correctly describes the language accepted by A?



FoodieQA: Procedural Single-Image QA

- The fine-grained taxonomy and careful labelling means that one can automatically create questions using templates

<dish>是哪个地区的特色菜？(What region is <dish> a specialty dish of?)
<dish>是哪个地区的特色美食？(In which region that <dish> is a local specialty?)
去哪个地方游玩时应该品尝当地的特色美食<dish>？ Which place should you visit to taste the local specialty food <dish>?

以下菜品是哪个地区的特色菜？

Which **region** is this food a specialty?



- ☒ A 江苏 (Jiangsu)
- ☐ B 京津 (Beijing & Tianjin)
- ☐ C 香港 (Hong Kong)
- ☐ D 广西 (Guangxi)

5. Binding: Degree of Multimodality

- The content expressed in textual data depends on the purpose

Social media platforms often form 'echo chambers' that encourage users to only read content that confirms beliefs they already hold (Getty)

Weak

(Mined)



A woman in a grey suit is giving a speech

Strong

(Crowdsourced)

Rewriting crawled text improves performance
on a variety of downstream multimodal tasks

COCO

$$P(x|v)$$

Distance(x, v)

- Images covering 80 common objects in context with multiple human-authored captions.
- Commonly used in image captioning or cross-modal retrieval
- Object segmentation data too!

some sheep walking in the middle of a road
a herd of sheep with green markings walking down the road
a herd of sheep walking down a street next to a lush green grass covered hillside.
sheared sheep on roadway taken from vehicle, with green hillside in background.
a flock of freshly sheered sheep in the road.



Conceptual Captions



- Used for pretraining
- 3/12M images released with *normalized* English captions.
- Normalization is not public.
- Due to *linkrot*, much less data is currently available.



VQAv2


- 1.1M image–question pairs with balanced distribution of answers
- Task with multimodal inputs:
 - Image
 - Question

Where is the child sitting?
fridge arms



Visual Commonsense Reasoning

- 290,000 multiple-choice VQA examples derived from movies with MCQA rationales



Why is [person4] pointing at [person1]?

- a) He is telling [person3] that [person1] ordered the pancakes.
- b) He just told a joke.
- c) He is feeling accusatory towards [person1].
- d) He is giving [person1] directions.

Rationale:


- a) [person1] has the pancakes in front of him.
- b) [person4] is taking everyone's order and asked for clarification.
- c) [person3] is looking at the pancakes both she and [person2] are smiling slightly.
- d) [person3] is delivering food to the table, and she might not know whose order is whose.

BBC-Oxford British Sign Language

$$P(y|x, v)$$

$$\text{Distance}(x, v)$$

- Sign spotting, sentence localization, and sign-language translation tasks
- 1,400 hours of signed shows
 - Factual, entertainment, drama, comedy, children's shows



MMMU-Pro

- 11.5K English multimodal questions from college exams, quizzes, and textbooks
- Collected by university students from online resources, adhering to copyright restrictions.

24 hoursanswers Subjects Online Tutoring Homework Help Homework Library Tutors Online Classes More - [Get Help Now](#)

Home > Homework Library >

No. 23: Each of the following situations relates to a different company. [image 1] For company D, find the missing amounts.

	Company A	Company B	Company C	Company D
1 Revenues	?	\$1,480,500	\$103,950	\$1,054,116
2 Expenses	\$455,490	1,518,300	78,120	?
3 Gains	0	?	4,725	8,505
4 Losses	32,760	0	5,670	39,312
5 Net Income or (Loss)	32,130	39,690	?	(58,275)


Options:

(A)\$1,081,584	(B)\$1,100,584
(C)\$1,034,325	(D)\$1,200,325
(E)\$1,125,325	(F)\$1,210,732
(G)\$1,150,732	(H)\$1,098,650
(I)\$1,075,732	(J)\$1,050,650


[View Available Computer Science Tutors](#)
624 tutors matched

VisionArena

- Diverse multilingual dataset with 238K human conversations with VLMs




In statistics, how do we interpret this distribution?



Diagram

Extract all text.




OCR

Who is this?



Entity Recognition



Many Many More

- Visual Storytelling, e.g. VIST
- Grounded Referring Expression, e.g. Flickr30K Entities, Visual Genome
- Visual Entailment, e.g. SNLI-VE
- Dense Image Captioning, e.g. Localized Narratives, DCI
- Image Captioning Evaluation Metrics, e.g. FOIL, Flickr8K-Expert
- Vision & Language Navigation, e.g. RxR
- Visual Common Sense Reasoning, e.g. VCR, WHOOPS!
- Text-to-Image Generation, e.g. DALLEval
- Abstract reasoning, e.g. CLEVR, KiloGram, CRAFT
- Sign Language Processing, e.g. How2Sign, SignNet-1M

- *and more and more and more and more*

Ethical Issues

- Multimodal datasets are increasingly scraped from the web with *unknown degrees of conformance*, or information about, licensing.



CC BY: This license allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, so long as attribution is given to the creator. The license allows for commercial use.

- By 2022 there were an estimated 2.5B CC-licensed objects online.
 - Current website now states that there are "tens of billions of CC-licensed works."

A Problem with Scale-First Thinking

- Scale lets you build systems that perpetuate harmful stereotypes



(Eileen Collins, American astronaut)

$\cos(v, x)$

0.276

“This is a portrait of an astronaut with the American flag”

0.308

“This is a photograph of a smiling housewife in an orange jumpsuit with the American flag”

Q: How can we collect multimodal data that better reflects the diversity of the world?

Vision-Language research is *English-first*

Languages

- Mostly in English
- Or some Indo-European Languages



ENG: An **unusual** looking vehicle ...

NLD: Een mobiel **draaiorgel** ...


Example from [van Miltenburg+ 2017](#)

Image sources


- Mostly from ImageNet or COCO
- Reflecting North American and European cultures

Implications for V&L models

- Narrow linguistic/cultural domain
- No way to assess their real-world comprehension




Density map of geographical distribution of images in ImageNet ([DeVries+, 2019](#))



FoodieQA: A Multimodal Dataset for Fine-Grained Understanding of Chinese Food Culture

Wenyan Li, Xinyu Zhang, Jiaang Li, Qiwei Peng, Raphael Tang, Li Zhou, Weijia Zhang,
Guimin Hu, Yifei Yuan, Anders Søgaard, Daniel Hershcovich, Desmond Elliott



香港中文大學
The Chinese University of Hong Kong



Motivation

- Existing multimodal datasets and models focus on food recognition, recipe generation, etc.
 - Regional difference in food culture is understudied
- **FoodieQA** is
 - manually curated
 - fine-grained
 - image-text dataset to evaluate food culture understanding across China



Sichuan




Guangdong

Tasks and Data Overview

Multi-Image VQA

哪一道菜属于川菜中的凉菜？ Which is a **cold dish** in **Sichuan cuisine**?



Single-Image VQA

以下菜品是哪个地区的特色菜？
Which **region** is this food a specialty?



- A** 江苏 (Jiangsu)
- B** 京津 (Beijing & Tianjin)
- C** 香港 (Hong Kong)
- D** 广西 (Guangxi)

Text QA

白切鸡的口味特色是？ What is the **flavor** of 白切鸡？

- A** 麻辣 (spicy)
- B** 松软 (soft)
- C** 外焦里嫩 (crispy-tender)
- D** 咸 (salty)

- FoodieQA comprises **three multiple-choice question-answering tasks**.
- Models answer questions based on **multiple images, a single image, or text-only descriptions**.

Collection and Annotation Process



1. Select and define cuisines


2. Define the fine-grained annotation taxonomy

3. Collect and label private images

4. Question construction

Multi-Image VQA

哪一道菜属于川菜中的凉菜? Which is a **cold dish** in **Sichuan cuisine**?



Single-Image VQA

以下菜品是哪个地区的特色菜? Which **region** is this food a specialty?



- ☒ A 江苏 (Jiangsu)
- ☐ B 京津 (Beijing & Tianjin)
- ☐ C 香港 (Hong Kong)
- ☐ D 广西 (Guangxi)

Text QA

白切鸡的口味特色是? What is the **flavor** of 白切鸡?

- ☒ A 麻辣 (spicy)
- ☐ B 松软 (soft)
- ☐ C 外焦里嫩 (crispy-tender)
- ☐ D 咸 (salty)

1. Define 14 Cuisine Types



2. Fine-grained Taxonomy



3. Private Image Collection

× 美食照片上传-Tencent Survey ...

*01 请上传您的美食照片

请选择尽量在正中的且清晰的美食照片上传。您如使用微信填写本问卷上传可能有延迟，建议使用浏览器打开（微信右上方）

+
添加图片/
文件

① 单个图片/文件大小不超过300MB

下一步

*02 您的美食属于哪个菜系/地区?

☐ 川菜（四川，重庆）

☐ 湘菜（湖南）

☐ 鲁菜（山东）

☐ 西北菜（陕西，甘肃等地）

Upload food image

Select cuisine type

× 美食照片上传-Tencent Survey ...

*03 美食名称

请尽量提供您所上传美食的准确名称，如**麻婆豆腐**而不是**豆腐**

10个汉字以内

Dish name

*04 您的美食来自

☐ 家里自制

☐ 餐馆

☐ 学校或单位食堂

☐ 外卖

Cook at home
Restaurant
Canteen
Delivery

*05 您确认所上传的图片**不是网络图片**
(即非网络下载图片，未上传至网络/社交媒体)，且允许我们将其用于非营利的科研工作。

Cook at home
Restaurant
Canteen
Delivery

*05 您确认所上传的图片**不是网络图片**
(即非网络下载图片，未上传至网络/社交媒体)，且允许我们将其用于非营利的科研工作。

我们收集的非网络图片将用于大数据模型的评测的相关研究，非网络图片对于我们的评测公正度非常重要。如您有任何疑问和请求，请联系：

☐ 确认

☐ 不同意

Confirm phot is not
uploaded/downloaded

提交


4. Annotation and Question-construction

- Human-written questions that require the image to answer; belong to one+ part of the taxonomy; and only one answer is correct.




Total Questions	403
Cohen's Kappa (κ)	0.83
Human Accuracy	0.81
Avg Question length	12.9
Multi-hop question	25%
# Question Types	14
# Unique Images	389

Multi-Image VQA Results



1. Big difference in performance of open-weight vs closed models.
2. Clear difference when the questions are posed in English vs Chinese.
3. Some variation in results for different prompt templates.

Image Source Matters



Publicly available images may have been seen during model pretraining

Take-away Messages

- FoodieQA is a multimodal dataset designed to evaluate fine-grained understanding of Chinese food culture through VQA and text QA
- Models are sensitive to the question language: bilingually trained Chinese–English models perform much better in Chinese
- Important for us to work with non-public images to avoid potential data distribution leakage, inflating model performance

Kaleidoscope: In-language Exams for Massively Multilingual Vision Evaluation

ICLR 2026



KØBENHAVNS
UNIVERSITET

Salazar and 44 collaborators


Global Multimodal Exams

Evaluating Multimodal LLMs Across Language Boundaries

Physics

Question #1925

त्रिज्या R एवं द्रव्यमान m के एक एकसमान खोखले बेलन के चारों तरफ एक इन्द्रविहीन डोरी से एक द्रव्यमान ' m ' अवलम्बित है। यदि डोरी बेलन पर फिसलती नहीं है, तब छोड़े जाने पर द्रव्यमान किस त्वरण से गिरेगा?



1) g

2) $\frac{2g}{3}$


3) $\frac{g}{2}$

4) $\frac{5g}{6}$

Geography

Question #11065

Kuriai geografinėi zonai (biomui) būdingas paveiksle pavaizduotas miškas?






1) Drėgnųjų atogrąžių (tropikų) miškų

2) Mišriųjų miškų


3) Plačialapių miškų

4) Spygliuočių miškų (taigos)



Dataset Properties

- 20,911 questions collected from a community of enthusiastic open-scientists
 - 18 languages
 - 14 subjects
 - 55% multimodal questions



Eight Types of Visual Information

Diagrams

High School Exam Physics

Arabic

Question: قراءة الأميتر الحراري في الدائرة الموضحة تساوي

Options:

- A. 0.2 A
- B. 2 A
- C. 0.02 A
- D. 20 A

Graphs

Physician Exam Language

Flemish

Question: Figuur 1A toont de A-weging voor het menselijk gehoor. Deze figuur leert ons dat

Options:

- A. de mens tonen rond de 1000 Hz het beste hoort.
- B. mensen tonen van 10.000 Hz niet meer kunnen horen.
- C. een toon met dezelfde fysieke geluidsterkte altijd even intens wordt gehoord.
- D. bij gelijke geluidsterkte, een mens 100 Hz zachter hoort dan 50 Hz.

Figures

HSC Exam Geography

Bengali

Question: উদ্দীপকের 'ক' ও 'খ' বাস্তুপ্রকারের সাধারণ বৈশিষ্ট্য- I. নদী-বন্দিন দিকে প্রবাহিত হয় II. ভাল দিকে বেঁকে যায় III. সার উচ্চতারিণি নিচের কোনটি সঠিক?

Options:

- A. i ও ii
- B. i ও iii
- C. ii ও iii
- D. i, ii ও iii

Formulae

SAT Mathematics

English

$4x^2 - 9 = (px + t)(px - t)$

Question: In the equation, p and t are constants. Which of the following could be the value of p?

Options:

- A. 2
- B. 3
- C. 4
- D. 9

Maps

Question: Діяльність якого гетьмана можна характеризувати, спираючись на подану карту?

Options:

- A. Б. Хмельницького
- B. І. Вигоцького
- C. Д. Многогрішного
- D. І. Саволюка

Photos

UP-CET Social Sciences

Hindi

Question: चित्र में दिए किले को पहचानिये।

Options:

- A. जोधपुर किला
- B. ग्वालियर किला
- C. लाल किला
- D. अमर किला

Tables

Portuguese

$\frac{p}{q+1}$	$\frac{p}{2q+1}$	$\frac{p}{3q+1}$	$\frac{p}{4q+1}$	$\frac{p}{5}$
-----------------	------------------	------------------	------------------	---------------

Question: Um segmento de reta está dividido em duas partes na proporção direta quando o todo está para uma das partes na mesma razão em que essa parte está para a outra. Essa constante de proporcionalidade é comumente representada pela letra grega φ , e seu valor é dado pela solução positiva da equação $\varphi^2 = \varphi + 1$. Assim como a potência φ^n , as potências superiores de φ podem ser expressas da forma $a\varphi + b$, em que a e b são inteiros positivos, como apresentado no quadro. A potência φ^7 , escrita na forma $a\varphi + b$ (a e b são inteiros positivos), é

Options:

- A. $7\varphi + 2$
- B. $9\varphi + 6$
- C. $11\varphi + 7$
- D. $13\varphi + 8$

Text



ABC ত্রিভুজ B কোণের পরিমাণ 80° এবং $AB=AC$ ।

Question: যদি E এবং F AB এবং AC-কে এমনভাবে ছেদ করে যেন $EF \parallel BC$ হয়, তাহলে

Options:

- A. 30°
- B. 50°
- C. 60°
- D. 70°

Cross-Cutting Evaluation




Clearly better text-only and Latin-script performance

Subject-type Analysis

- Clear difference in performance when we break the results down into **HSS** and **STEM**
 - What does this tell us about the information that can be *easily learned* from pretraining and instruction tuning?


	Closed Weights		
	Gemini	Claude	GPT-4o
<i>Humanities & Social Sciences</i>			
Economics	64.1	63.8	66.7
Geography	72.8	81.5	80.4
History	78.7	83.7	86.4
Language	83.5	85.5	85.8
Social Sciences	85.7	82.9	88.1
Sociology	92.3	93.4	93.2
<i>STEM</i>			
Biology	60.3	62.9	63.9
Chemistry	60.4	59.7	52.9
Engineering	57.3	64.4	56.3
Mathematics	48.6	44.4	44.0
Physics	57.8	58.7	54.7
<i>Reasoning, Health Science, and Practical Skills</i>			
Reasoning	52.0	53.3	51.0
Medicine	70.2	73.8	75.6
Driving License	64.4	64.2	73.1

Image-Type Analysis



Models struggle with diagrams, figures, and formulae

Qwen2.5-VL Scaling Performance



Promising log(linear) performance improvement in model scale

Take-away Messages


- Large-scale open-science initiative can yield good data
- Clear differences in performance on several axes:
 - Language
 - Closed vs open-weights models
 - Image type matters a lot
 - HSS vs STEM
- Open-weights models need improvements for all image types
- Performance scaling suggests promising avenues for development!

2. Data Representation

Three Levels of Representation

- Perceptual (i.e., language oriented)
- Pre-processed features (i.e., object/stuff oriented)
- Raw inputs

- ☐ Yellow
- ☐ Has wheels
- ☐ Metal
- ☐ Five-door
- ☐ Can transport
- ☐ ...



Perceptual Representations

Key Idea: Describe perception using (human-given) symbolic/linguistic features.

- Psychology studies asking people to write down words that are triggered by stimuli.
- Stimuli: 541 noun concepts
- Perceptual norms can be categorized into the likely knowledge source

Moose		
is large	27	visual-form and surface
has antlers	23	visual-form and surface
has legs	14	visual-form and surface
has four legs	12	visual-form and surface
has fur	7	visual-form and surface
has hair	5	visual-form and surface
has hooves	5	visual-form and surface
is brown	10	visual-color
hunted by people	17	function
eaten as meat	5	function
lives in woods	14	encyclopedic
lives in wilderness	8	encyclopedic
an animal	17	taxonomic
a mammal	9	taxonomic
an herbivore	8	taxonomic

Perceptual Norms: Pros / Cons

Pros


- Seemingly simple task
- Relatively rich features

Cons

- Can it scale?
- Handling ambiguity

How much can *deep understanding* be made from language alone?

What does it mean to only understand symbols as defined by other symbols?




Searle. (1980). Minds, Brains and Programs. *Behavioral and Brain Sciences*, 3: 417–57

Ramos et al. (2023) LMCap: Few-shot Multilingual Image Captioning by Retrieval Augmented Language Model Prompting. ACL Findings

Zheng et al. (2023) Socratic Models: Composing Zero-Shot Multimodal Reasoning with Language. ICLR


Spatial and Pooled Visual Features

- Earliest work in neural-network era used pooled or spatial preserving features from a pretrained Convolutional Neural Network.



Pre-processed Visual Features

- Faster R-CNN region-based feature vectors
 - Trained on the Visual Genome Dataset
 - The Region Proposal Network suggests the location of *regions of interest*.
 - RoI pooling performs spatial pooling in the final CNN layer to give a 2048D vector.



Pre-processed: Pros / Cons

Pros

- Long-established practice
- Usually an offline process: do it once and forget

Cons


- Large datasets require specialized storage
- Not obvious how to randomly augment data
- Specialist knowledge can be opaque to newcomers

Raw Input


- Directly process data from the raw images or speech signal.
- Images:
 - Vision Transformer (ViT)
 - Swin Transformer
- Speech
 - Spectrogram Transformer
 - AudioMAE

Vision Transformer

- Good news! You are already almost an expert in how a standard Vision Transformer model works!
 - Split image into K patches
 - Embed each patch
 - Add position information
 - Encode using Transformer blocks that include an **extra pre-norm layer** for stability.



Nomenclature and Patch Count




Extracting ViT Features

- Extract **pooled features** or **patch-level features**


To extract visual information from an image x^i , we use the visual encoder of a pre-trained CLIP [29] model. Next, we

For the CLIP encoders, we extract the feature grid before the pooling layers, resulting in an $N \times N$ grid, where $N = 7, 7, 12$ for the ViT-B/32, RN50x4 and RN50x16 variants of CLIP respectively.



Multiple and High-Resolution Visual Inputs

- Multiple ideas in the recent literature:
 - Independent processing of different images or image patches
 - Compressing the sequences of patches
 - Similarity-based (*grouping*)
 - Sparse attention
 - Query-based
 - Multi-step processing with region (dynamic) selection and cropping
 - Combinations of different approaches




Scaling on Scale

- Improve performance by combining original and higher-resolution processing without additional pre-training or finetuning




Process High-Resolution Inputs

- Learn to combine multi-scale inputs (e.g., with local and/or global LoRA adapters) from vision encoders with language models



Compressing Visual Tokens



Shao et al., A Survey of Token Compression for Efficient Multimodal Large Language Models. TMLR 2026

Yang et al., VisionZip: Longer Is Better but Not Necessary in Vision Language Models, CVPR 2025

Shi et al., Attend Before Attention: Efficient and Scalable Video Understanding via Autoregressive Gazing, CVPR 2026

Raw input: Pros / Cons

Pros

- Data augmentation is straightforward because you always have the raw input
- Fewer preprocessing steps means fewer creeping errors

Cons

- Smaller batches with an extra model on the GPU
- Potentially many input tokens, and we need strategies to handle this


Summary

- Three main options for how to represent your multimodal inputs
 - Language-oriented
 - Object / stuff oriented
 - Raw inputs
- **Raw inputs** are the current favoured approach to visual encoding because you can update the weights in the model


3. Modelling

Main Approaches


Cross encoding



Cross-Attention




Visual Prefix




Cross-encoding Models

- Emerged as a key modelling approach in 2019 with many concurrent methods for creating visually-grounded BERT models.
- This is a form of *model-based fusion*
- The backbone consists of two components:
 - language encoder
 - visual encoder






Single- & Dual-Stream Architectures

- Single-stream
 - Concatenate inputs into one sequence
- Dual-stream
 - Process modalities independently
 - Intra-modal
 - Inter-modal




ViLBERT

- Dual-stream model
- Initialized from BERT
- Visual features extracted from 10-36 regions using Faster-RCNN
- Pretrained on Conceptual Captions
 - Masked Language Modelling
 - Masked Region Classification
 - Image-Text Matching




ViLT

- Single-stream model
- **Initialized from BERT**
- Visual features extracted from ViT-B/32
- Pretrained on Conceptual Captions, Visual Genome, COCO, SBU Captions
 - Masked Language Modelling
 - Image-Text Matching
 - Word-Patch Alignment




FLAVA



- Dual-stream Visual features extracted from ViT-B/16
- Pretrained on PMD70M
 - Masked Language Modelling, Masking Image Modelling
 - Image-Text Matching, Masked Multimodal Modelling
 - Global Contrastive Matching

Dual-encoding Models

- Emerged as a sample-efficient alternative to cross-encoding
- The backbone consists of two separate components:
 - language encoder
 - visual encoder



CLIP

Language

The red horse


Vision



- 12 Layer Transformer Encoder
- ViT or ResNet Visual Encoder
- Maximize the similarity of the embeddings of paired examples (I, T):

$$\mathcal{L}_{\text{InfoNCE}} = -\mathbb{E} \left[\log \frac{f(\mathbf{t}, \mathbf{i})}{\sum_{\mathbf{t}' \in T} f(\mathbf{t}', \mathbf{i})} \right]$$

- Large pretraining dataset of unclear provenance




Cross-Attention

Attend

Language Model




The red horse




Visual Prefix Learning

- Main goal: force the visual representations into an LLM-compatible space through a learned projection



Current Strategies for Training VLMs



- Three main parts
 1. Pretrained models: vision encoder & language model
 2. Define a connection layer between the models
 3. Train the connection with multilingual multimodal data
 - a. Continued pretraining and instruction tuning

Modality-Specific Components

	Vision Encoder	Language Model
LLAVA	CLIP ViT-L/14	Vicuna-13B
Qwen-VL	OpenCLIP ViT-bigG	Qwen-7B
MM1	ViT-L	1.3B LLM
PaliGemma	SigLIP-So400M/14	Gemma-2B

<https://ai.google.dev/gemma/docs/paligemma>

<https://llava-vl.github.io/>


<https://github.com/QwenLM/Qwen-VL>

Connector Design

- PaliGemma: Map the image output embeddings into the language model word embedding space.
- Qwen-VL:

Position-aware Vision-Language Adapter: To alleviate the efficiency issues arising from long image feature sequences, Qwen-VL introduces a vision-language adapter that compresses the image features. This adapter comprises a single-layer cross-attention module initialized randomly. The module uses a group of trainable vectors (Embeddings) as query vectors and the image features from the visual encoder as keys for cross-attention operations. This mechanism compresses the visual feature sequence to a fixed length of 256. The ablation about the number of queries is shown in Appendix E.2. Additionally, considering the significance

- MM1: Convolutional-Abstractor
 - ResNet Block followed by an Adaptive Pooler



Training Datasets

- LLAVA: 595K image–caption examples filtered from CC3M

Qwen-VL 1.4 billion examples (77% English / 23% Chinese)


MM1 2+ billion mixture of image–text examples

PaliGemma 1 billion mixture of multilingual image caption, VQA, and in-the-wild datasets

- The larger models are pretrained on **in-house data**
 - PaliGemma: WebLI (1B+), Qwen-VL (220M), MM1 (1B+)



Training Strategy

- VLMs are trained using multi-stage fine-tuning strategies with different types of data and different image resolutions





Trends in Smaller Models



- How quickly will we realize these benefits in smaller models?





Balance Models and Increase Context



Learned Position Embeddings Help!




Chain-of-Thought Data Hurts



Molmo and Molmo 2

- Two recent open-weights and open data VLMs that are very transparent about the data used for training and what mattered to drive model performance
 - Design brings together many topics previously discussed in this lecture



PixMo Datasets


- Collect highly detailed descriptions through voice instead of typing

PixMo


Captions



AskModelAnything



Pointing




Synthetic



Molmo


Fine-grained Understanding



User Interaction



Pointing and Counting



Visual Skills



Molmo Fine-tuning Data Mixture

- PixMo data and existing datasets
 - Annotated by humans
 - Synthetically generated
- Clear documentation about pretraining and fine-tuning hardware (H100s)

		fine-tune	
	GPUs	time	GPU hr.
1B-E	64	13.3	850
7B-D	128	11.2	1.4k
7B-O	128	13.5	1.7k
72B	256	32.4	8.3k



Molmo: What drives performance?

- More overlapping image crops

cropping	cap F_1	11-avg
single	46.7	62.8
multi, no overlap	53.4	75.7
multi, overlap	54.1	76.9

- Dropout in the LLM

pre-train, fine-tune	cap F_1	11-avg
off, off	53.1	74.6
off, on	53.1	76.6
on, on	53.7	77.0
on (text only), on	54.1	76.9


- Specialized fine-tuning data

data	11-avg
academic only	72.5
plus PixMo-Docs	74.0
PixMo-★ plus academic	76.9
remove PixMo-AMA	76.8
remove PixMo-CapQA	77.0
remove PixMo-Docs	75.8
remove PixMo-Clocks	76.9
remove pointing task	76.2

- Cleaned pretrained data

data	cap F_1	11-avg
stage 0.5 LAION	53.9	76.9
ShareGPT4V+o (158k images)	36.3	74.9
<i>PixMo-Cap images:</i>		
our raw transcripts only	45.2	76.4
our cleaned transcripts only	53.0	76.5
our raw & cleaned transcripts	54.1	76.9
captioned by GPT-4o	52.9	77.5

Molmo2 Data Mixture and Performance



Model	Short QA avg.	Long QA avg.	Average
API call only			
GPT-5 [114]	73.1	76.3	70.6
GPT-5 mini [114]	66.8	69.8	65.0
Gemini 3 Pro [45]	71.0	78.8	70.0
Gemini 2.5 Pro [25]	71.1	80.4	71.2
Gemini 2.5 Flash [25]	67.0	74.5	66.7
Claude Sonnet 4.5 [5]	62.8	66.4	59.6
InternVL3.5-4B [149]	62.0	56.5	53.4
InternVL3.5-8B [149]	63.0	57.1	54.1
Qwen3-VL-4B [169]	63.7	62.7	58.1
Qwen3-VL-8B [169]	65.3	63.5	59.5
Molmo2-4B	<u>69.3</u>	<u>64.5</u>	<u>62.8</u>
Molmo2-8B	69.9	64.1	63.1
Molmo2-O-7B	68.1	59.2	59.7

Summary

- Cross-encoding:
 - Many advances in which parts of the input contribute to loss
 - Shift from regions-of-interest to Vision Transformers
- Dual-encoding:
 - Excellent cross-domain transfer to a wide range of problems
- Visual Prefix Learning:
 - Exploit the benefits of single-modality pretraining

Q: How can we squeeze more performance out of Small VLMs?

Efficient Test-Time Scaling for Small Vision-Language Models

ICLR 2026



M. Onurcan
Kaya




D. Elliott



Dim. P.
Papadopoulos

Test-Time Scaling (TTS)

- Rather than relying on ever-larger pretraining budgets, test-time methods use dynamic inference strategies that allow models to “think longer” on harder problems.



Source: <https://blogs.nvidia.com/blog/ai-scaling-laws>

What to Scale in TTS?

- Lengthen a single reasoning trajectory:
 - Chain-of-Thought ([Wei et al., 2022](#))
 - OpenAI o1 ([Jaech et al., 2024](#))
 - DeepSeek R1 ([Guo et al., 2025](#))
- Increase the number of trajectories and aggregate them:
 - Self-Consistency ([Wang et al., 2023](#))
 - Best-of-N ([Brown et al., 2024](#); [Irvine et al., 2023](#))
 - REBASE ([Wu et al., 2024](#))

Motivation

- Small VLMs are efficient and accessible, but they have weaker generalization abilities and downstream task performance.




Question: Is the pigeon standing in this image?

SmoIVLM2-2.2B Answer:

No, the pigeon is perched on top of the coin return machine.


Key Problem

- Existing test-time scaling methods require heavy external verifiers or large compute budget; not feasible for small-models or low-compute scenarios.




The figure is adapted from “Large Language Monkeys: Scaling Inference Compute with Repeated Sampling”.

Overview of Our Framework



Augmentation Examples



Who is the author of this book?

**Image
Transforms**




**Text
Transforms**

Who is the author of this book?
In other words, Who is the author of this book?

Results

- Nine diverse tasks:
- **Visual question answering:** ChartQA, OCRBench, OCRVQA, GQA, TextVQA
- **Multiple-choice questions:** AI2D and MME-RealWorld
- **Yes/no questions:** AMBER
- **Image captioning:** COCO Captions



Lots of Experiments

How to augment?

Optimal method?

Why augment?

Where to aggregate?


Optimal layer?

How to adapt?


How to aggregate?

Optimal granularity?


Optimal method?



Diversity-Inducement Methods





- **Temperature Sampling** introduces diversity by sampling from a softened probability distribution.
- **Input Perturbations** applies classic semantic-preserving augmentations to inputs and then **decodes greedily**.
- Evaluate using two TTS strategies:
 - **Self-Consistency** (majority voting)
 - **Self-Selector** (using the VLM itself)




Takeaway: Input Perturbations with greedy decoding performs best

Aggregation Levels




Takeaway: Token-level captures fine-grained confidence signals.

Aggregation in Early Layers




- What is the optimal point for feature aggregation within the model architecture?




Takeaway: Language-heavy tasks benefit from late-layers, vision-heavy tasks are best with early-layer aggregation.

Number of Augmentations




- We study how performance scales with the number of augmented inputs to understand the optimal balance between computational cost and accuracy.




Takeaway: Performance gains plateau after ~16 augmentations


Computational Overhead



Peak GPU Memory (GB)




Runtime (sec)




Takeaway: Parallel implementation minimizes latency but requires substantial memory, while sequential implementation conserves memory at the cost of increased runtime.

Text Augmentations




- **Image-only:** Control experiment.
- **AugGPT** uses ChatGPT to paraphrase.
- **Self-paraphrasing** uses the LLM of the VLM to paraphrase the input prompt.
- **Classical augmentation:** Fast and simple. Keyboard errors, word splitting, deletion, and sentence reordering.
- **Consistency enforcement:** Append the original prompt after each augmentation:

Augmented + "In other words, " + Original.



Takeaway: Consistency enforcement matters

Test-Time Adaptations




- ① **Aggregation weights optimization:**
 - Learns adaptive aggregation weights (instead of simple uniform averaging).
 - Minimizes marginal entropy.
 - Low overhead, efficient.
- ② **Model parameter adaptation:**
 - Iterative pseudolabeling + fine-tuning.


	Baseline	TTAug	①	②
ChartQA	74.2	76.1	76.1	76.7
OCRBench	72.9	73.7	73.0	70.5
OCRVQA	0.0	12.6	11.9	13.8
GQA	0.0	5.5	5.2	13.5
TextVQA	73.2	72.4	74.2	70.5
AI2D	68.5	69.6	69.7	67.4
MME-RW	27.8	31.9	30.9	31.4
AMBER	68.7	75.9	76.9	72.8
COCO	9.1	16.9	16.4	35.9
Mean	43.8	48.3	48.3	50.3

Takeaway: Model parameter adaptation gives highest performance gains.

Cross-Model Generalization



- We apply the best SmoVLM2-2.2B configuration to other architectures and parameter scales.



Takeaway: Our methods provide improvements across diverse models, though transferability varies by family and size.

Take-away Messages

- Introduced two efficient test-time scaling methods: **Test-Time Augmentation** and **Test-Time Adaptation**
- Improved accuracy across nine diverse benchmarks with minimal overhead.
- **Key Insight:** Token-level aggregation and augmentation-based diversity inducement outperform traditional approaches.


4. Understanding Multimodal Models

Beyond Task Performance

- Many questions about what drives the success of these models?
 - Better contextualization: make better use of the multimodal inputs
 - Acquire certain “skills”, e.g. counting or localization
 - Understand linguistic structures
 - Something else?
- Model-internal behaviour
 - Attention mechanism patterns
 - Mechanistic interpretability (emerging)
- Probing
 - Tasks related to different skills

FOIL Captions

- Do V&L models really understand the relationship between words and images?
- Crowdsourced datasets that contain contextually plausible but incorrect image–text pairs




Vision and Language Understanding Evaluation

- Suite of five model probing tasks
- **Modality Influence:** Estimate the layer-wise contribution of each modality to the [CLS] embedding:

$$I_{M,j} = \sum_{i \in S} \mathbb{1}(i \in M) \cdot \alpha_{ij}$$

- The UNITER model relies more on textual features when fusing modalities throughout the model



VALSE Benchmark

- Test visio-linguistic capabilities with image-sentence foil pairs
- Image-sentence matching task
 - Existential quantifiers
 - Semantic number
 - Counting
 - Prepositional relations
 - Action replacement / swap
 - Co-reference




A small copper vase
with **some flowers** /
exactly one flower in it.

Metric	Model	Avg.
	Random	50.0
acc_r	GPT1*	60.7
	GPT2*	60.1
	CLIP	64.0
	LXMERT	59.6
	ViLBERT	63.7
	12-in-1	75.1
	VisualBERT	<u>46.4</u>


$$p(\text{caption}, \text{img}) > p(\text{foil}, \text{img})$$

VL-CheckList


- Evaluate V&L models based on automatic manipulations to vision and language data.
- Image-Sentence matching task
- Radar chart overviews based on object / attribute / relationship variations



(a) Language Variation



(b) Vision Variation



Subject-Verb-Object Probes

- Large-scale dataset with SVO triplets mined from Conceptual Captions and 14K images and with crowdsourced captions
- Foil detection formulation



WinoGround

- 1,600 text-image pairs to evaluate compositional understanding



some plants
surrounding a
lightbulb




a lightbulb
surrounding
some plants

- Images sourced **with permission** from Getty.
- Differences are categorised into: swap dependent, swap-independent, and visual differences

Understanding

- Train linear probes to predict the semantic attributes of concepts from frozen encoder representations
- Hard generalization task evaluating on *unseen* concepts
 - Train: **Cat**, **Dog**, **Cow** → `has_four_legs`
 - Test: **Table** → `has_four_legs`




Object Grounding and Hallucination

- Explicitly training a VLM to attend to aligned object regions during post-training does not seem to reduce hallucination




A **white** hound and a **cat** looking at the camera

"A white **hound** and a **cat** looking at the camera"




CHAIR

"A white hound and a cat looking at the camera"




FaithScore

Modality Dominance




- Pair of diagnostic evaluations that can be applied to any model that makes MLM and MRC predictions.





How do LLMs integrate vision?

- Some evidence of *direct* information extraction (Neo et al. 2025)




Understanding the Linear Projector

- The projected image representations are not a good fit for a language model sparse autoencoder until the final layers of the LLM



How Does Visual Information Flow?

- VLMs integrate visual information at different stages when responding to an input prompt (Zhang et al. CVPR 2025)






Legend:

- Question $\rightarrow$ Last (Green line)
- Image $\rightarrow$ Last (Yellow line)
- Last $\rightarrow$ Last (Red line)
- Image $\rightarrow$ Question (Blue line)
- Related Image Patches $\rightarrow$ Question (Green line)
- Other Image Patches $\rightarrow$ Question (Pink line)

Is visual information retained in LLMs?



- On vision-heavy tasks, the projection seems to lose critical visual features

Task (e.g., correspondence)



Understanding Reasoning Dynamics

- How do VLMs revise their predictions when they have multiple Chain-of-Thought steps to “think” about the problem?
 - Whether models are correct or incorrect, there is equal probability for their response
 - Behaviour is easily influenced by misleading interventions




Q: How are visual inputs represented in Vision-Language Models?

LATENTLENS: Revealing Highly Interpretable Visual Tokens in LLMs

ICML 2026



B. Krojer



S. Nayak



O. Mañas



V. Adlakha



D. Elliott




S. Reddy



M. Mosbach

New method for interpretation



Prior Work

Logitlens

unembedding matrix $\rightarrow$ NNs

dog
the
?!

LLM



dog the ?!

classes

“EmbeddingLens”

embedding matrix $\rightarrow$ NNs

Probing


Define classes $\rightarrow$ train

Issues


Insights about relation to LLM embeddings? **✗**

Need to train+define probe classes **✗**

LATENT LENS




LATENTLENS: details




GPT-5 validated
on 1,000 human
annotations
(Cohen's $\kappa = 0.68$)

Controlled Experimental Setup




Peeking inside VLMs



Finding 1:

EmbeddingLens shows medium/high interpretability even at input layer **for some models**


Peeking inside VLMs



Finding 2:

LogitLens shows high interpretability, but only at **late layers**, and only **some models**


Peeking inside VLMs



Finding 3:

With the *LATENTLENS*, vision tokens are much more interpretable **across all layers** and **all tested models**


Peeking inside VLMs



Insights also hold in several ablations, when we:

1. Apply to an off-the-shelf model (Qwen2-VL-7B-Instruct) ✓
2. **update** LLM during training → even more interpretability ↑
3. replace MLP with **linear mapping** → ~same interpretability ✓

Qualitative Examples



Additional Analysis

1. Phrases provide fine-grained interpretability
2. Visual text is most reliably interpretable




“large stone tower
with gold **clocks**”

(*OLMo+CLIP, L16*)


Logitlens:

"circular", ..., "bsites"

CLIP-ViT




DINOv2



Non-interpretable tokens?

- There are also non-interpretable tokens with low cosine similarity
- This raises the question: are they useful to the model downstream?
 - Yes!




Takeaways

- *LATENTLENS* provides interpretable textual descriptions of the visual input representations in VLMs
- Our findings generalize across combinations of vision encoders x language decoders *and* to pretrained models

5. Future Directions

Physical Understanding

- Predicting and explaining physical actions in the world will become of increasing importance as we create embodied agents





Q: How many objects are prevented by the tiny green triangle from falling into the basket?

Q: What is the color of the last object that collided with the tiny red circle?

Q: If any of the other objects are removed, will the tiny green circle end up in the basket?


Text-based Video Games

- Learning to act in procedurally-generated video game environments with rich contexts, action spaces, and long-term rewards




Multimodal Agents

- OmniAct combines multimodal understanding with program execution to solve a variety of tasks that humans perform with their computers





Multilinguality

- The majority of Vision and Language research is in English
- We need resources, models, and evaluations to create useful multilingual multimodal models
- High-quality data requires:
 - time
 - money
 - community engagement



Multimodality for Social Good

- How can we understand and limit the harm that our Vision-Language models will unleash on society?



Model	Type of Response	%	%
xGen-MM		14.0	54.0
Qwen-2-VL		7.3	53.0
MiniCPM-2.6		7.3	9.0
InternVL-2		5.8	12.8
Idefics-3		4.5	42.0
InternLM		2.8	15.3
Cambrian		2.5	13.8
GPT-4o		1.0	5.5
Gemini-1.5		0.3	7.3
Claude-3.5		0	2.5


Wrap-up

1. Datasets




some sheep walking in the middle of a road
 a herd of sheep with green markings walking down the road
 a herd of sheep walking down a street next to a lush green grass covered hillside.
 sheared sheep on roadway taken from vehicle, with green hillside in background.
 a flock of freshly sheered sheep in the road.



2. Representation



3. Modelling



4. Understanding

5. New Directions

Acknowledgements



E. Bugliarello



R. Ramos



A. Kádár



L. Specia



L. Barrault




D. Sánchez



N. Collier



G. Chrupała



W. Li



D. Oneață




S. Frank



B. Krojer



M. Fadaee



M. O. Kaya



B. Martins



E. Hasler



E. Ponti



S. Hooker



S. Reddy



L. de Grazia



I. Salazar