

ML4H 2021 Research Roundtable

Multimodal learning in healthcare and
representation learning on clinical data




Senior Chairs: Marzyeh Ghassemi, Polina Golland

Junior Chairs: Sneha Jha, Rishab Khincha

Clinical Data is inherently multi-modal

- Structured tabular data
- Time series
- Textual clinical notes
- Imaging data
- Histology slides
- Omics data
- Specific tests/analyses
- Metadata (age, gender, demographic information)



Comparison:

None.

Indication:

Chest pain, feels out of it.

Findings:

The Cardiomeastinal silhouette and pulmonary vasculature are within normal limits in size. The lungs are clear of focal airspace disease, pneumothorax, or pleural effusion. There are no acute bony findings.

Impression:

No acute cardiopulmonary findings.

Why is multi-modal data useful?

- Useful information is known to exist in an alternate modality
- Latent information in an existing modality
- Complementary information exists across modalities that can be exploited jointly

What directions in machine learning are well-suited to learning from multimodal data?

How to do multi-modal


- Stack up modalities
- Multi-modal ensembles
- Late fusion
- Representation learning: image-based text generation, combining representations for downstream tasks, contrastive learning

- Representation Learning via Maximization of Local Mutual Information: what are its key differences and what is the importance of the local features used?

Challenges posed by multi-modal data

Does multimodal learning present any fresh challenges in terms of privacy and interpretability?

Avenues for better interpretability methods by combining explanations from different modalities in a meaningful way?




(1) A mass is present in the superior segment of the left lower lobe and therefore malignancy must be considered. (2) Elsewhere, the left lung appears clear. (3) There is no pleural effusion. (4) Calcified pleural plaque is present in the right mid zone. (5) The right lung appears clear.

Representation learning X healthcare

Can representation learning change the way current models are trained in healthcare?

Massive pre-trained set of networks like ImageNet?

Something like CLIP for zero shot transfer?



Thank you!

Sneha Jha sneha.jha@imperial.ac.uk
Rishab Khincha rkhincha@utexas.edu