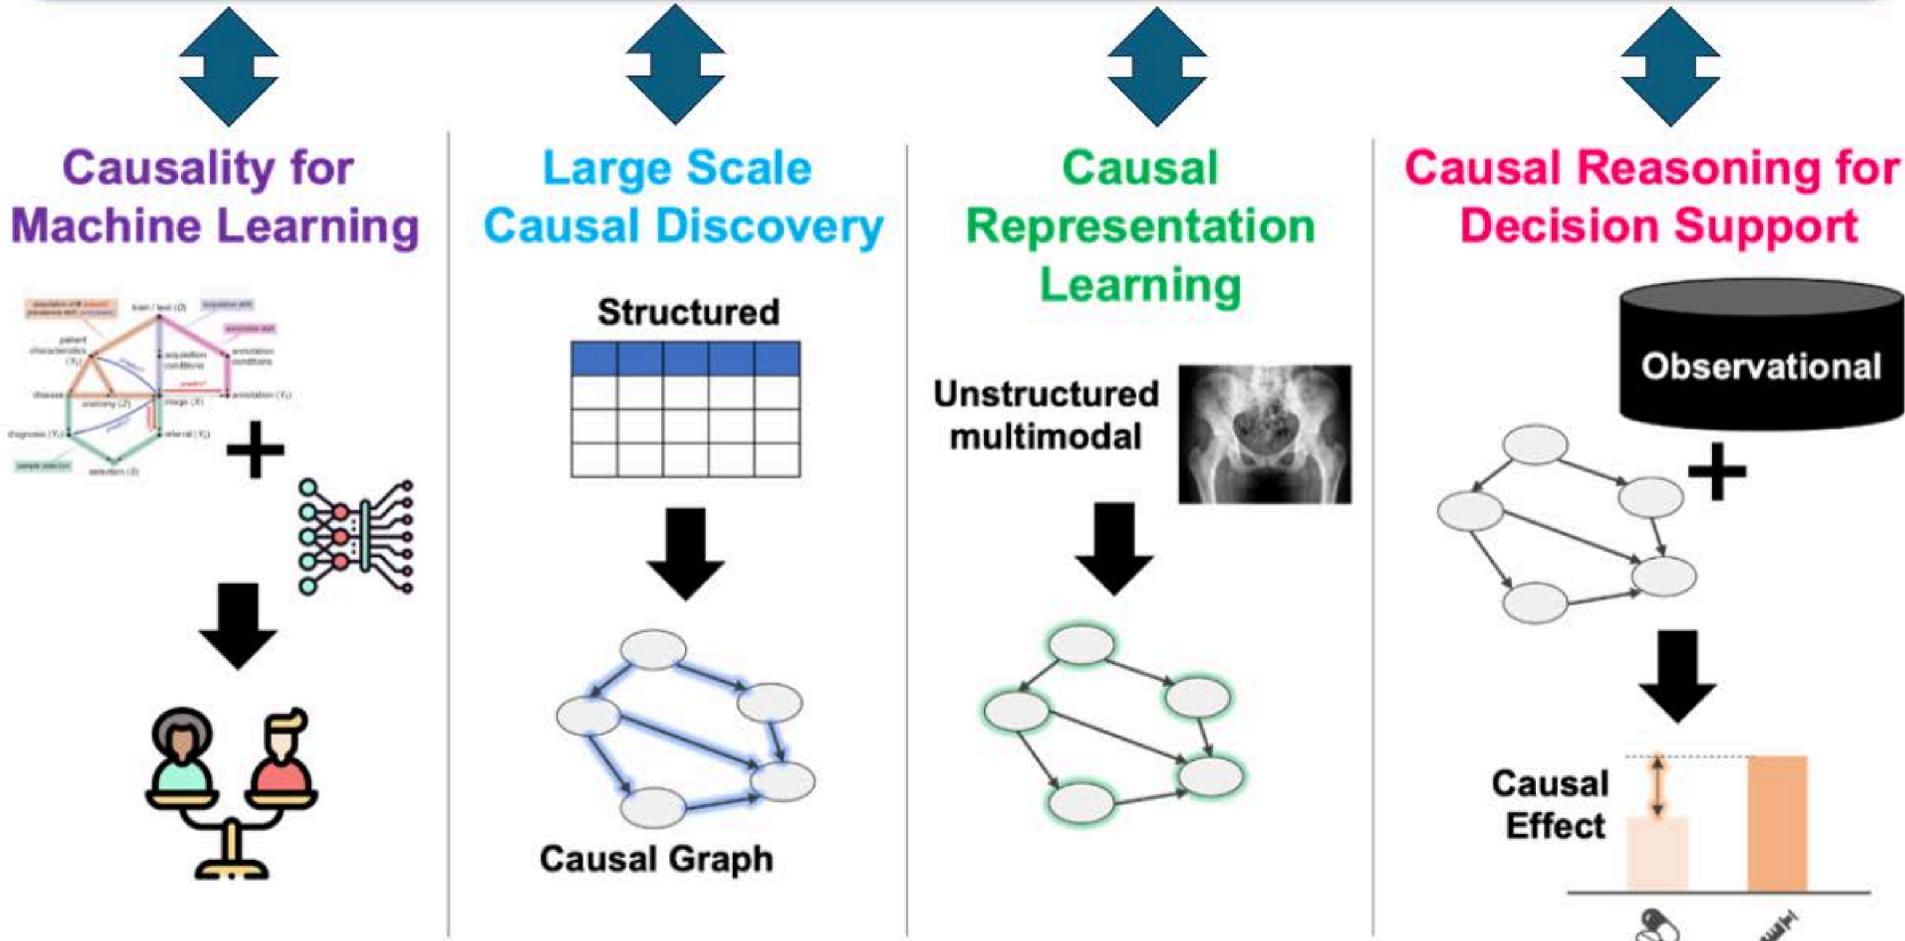




### Complex Real Data

federated, heterogeneous, noisy, sparse, and multimodal



### Causality for ML

Using causal ideas to do **better** machine learning  
 Can **improve** how we collect, annotate, and debias data and design fair algorithms  
 Characterisation of **failures** of ML predictive clinical models due to unaccounted selection biases  
 Exploring fine-tuning and other training regimes to improve out-of-distribution **generalisation**

### Large Scale Causal Discovery

Learning the **causal relationships** between variables in large dimensions  
 For example, can learn relationships in large **gene** networks  
 Using diffusion models and modern generative AI for **efficient search** of causal structures  
 Data-efficient **active loop** of causal discovery, experimentation and knowledge elicitation

### Causal Representation Learning

Learning to extract causally-related variables from **unstructured** data (e.g. images)  
 For example, can learn the **influence** between disease and anatomy in medical images  
 Explanation methods and **counterfactual** generation: algorithms, applications & benchmarks  
**Linking** low-level variables to high-level causal factors in domains such as brain disease progression

### Causal Reasoning for Decision Support

Learning how to make decisions **combining** observational data and causal priors  
 Can improve how we combine **evidence** of trials and observations to make treatment suggestions  
 Learning causal and predictive **behaviour** on extrapolation regions by combining data sources  
 Improving **sequential** actions: longitudinal discovery, treatment combinations and beyond