Probablistic Learning To Defer

Cuong Nguyen[1] Thanh-Toan Do[2] Gustavo Carneiro[1]

April 25th, 2025

[1]Centre for Vision, Speech and Signal Processing, University of Surrey, UK

[2]Department of Data Science and AI, Monash University, Australia

Introduction

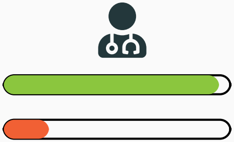

Reliability

Efficiency

Introduction

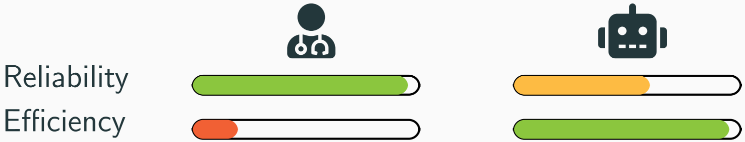

Introduction

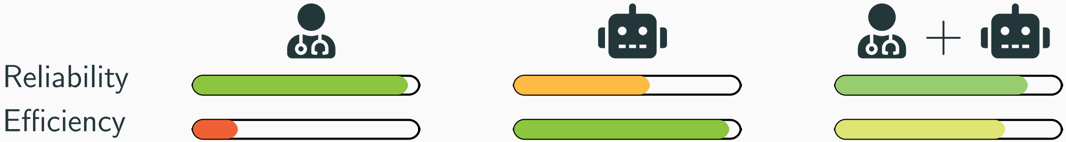

Learning to defer (L2D) aims to leverage:

- high *reliability* of human, and
- high *efficiency* of machine learning models.

Background - learning to defer

Background - learning to defer

Modelling

L2D is a mixture of:

Background - learning to defer

Modelling

L2D is a mixture of:

- M human experts, and

Background - learning to defer

Modelling

L2D is a mixture of:

- M human experts, and
- a machine learning model (i.e., classifier).

Background - learning to defer

Modelling

L2D is a mixture of:

- M human experts, and
- a machine learning model (i.e., classifier).

$\rightarrow$ latent variable model: expert selection $\mathbf{z}$

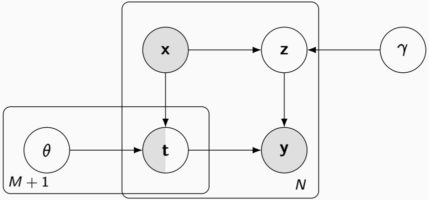

Figure 1: Graphical model of L2D.

Background - learning to defer

Modelling

L2D is a mixture of:

- M human experts, and
- a machine learning model (i.e., classifier).

$\rightarrow$ latent variable model: expert selection $\mathbf{z}$

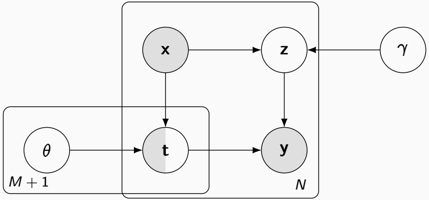

Figure 1: Graphical model of L2D.

Objective

$$\boxed{\max_{\gamma, \theta_{M+1}} \frac{1}{N} \sum_{i=1}^{N} \ln \Pr\left(\mathbf{y}_i, \prod_{m=1}^{M} \mathbf{t}_i^{(m)} \middle| \mathbf{x}_i, \gamma, \prod_{m=1}^{M+1} \theta_m \right).} \tag{1}$$

Background - learning to defer

Modelling

L2D is a mixture of:

- M human experts, and
- a machine learning model (i.e., classifier).

$\rightarrow$ latent variable model: expert selection $\mathbf{z}$

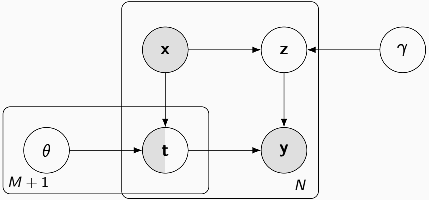

Figure 1: Graphical model of L2D.

Objective

$$\boxed{\max_{\gamma, \theta_{M+1}} \frac{1}{N} \sum_{i=1}^{N} \ln \Pr\left(\mathbf{y}_i, \prod_{m=1}^{M} \mathbf{t}_i^{(m)} \middle| \mathbf{x}_i, \gamma, \prod_{m=1}^{M+1} \theta_m \right).} \tag{1}$$

Learning is performed via the *Expectation - Maximisation* algorithm

Motivation

Learning to defer has limitations

Motivation

Learning to defer has limitations

⚠ requires *all* human experts must annotate every training sample

Motivation

Learning to defer has limitations

⚠ requires *all* human experts must annotate every training sample

- 🚫 impractical (e.g., each sample is annotated by few human experts), and

Motivation

Learning to defer has limitations

- requires *all* human experts must annotate every training sample
  - impractical (e.g., each sample is annotated by few human experts), and
  - costly, time-consuming, and even infeasible (e.g., radiology),

Motivation

Learning to defer has limitations

- requires *all* human experts must annotate every training sample
  - impractical (e.g., each sample is annotated by few human experts), and
  - costly, time-consuming, and even infeasible (e.g., radiology),
- most likely selects the *best human expert* all the time

Motivation

Learning to defer has limitations

- requires *all* human experts must annotate every training sample
  - impractical (e.g., each sample is annotated by few human experts), and
  - costly, time-consuming, and even infeasible (e.g., radiology),
- most likely selects the *best human expert* all the time
  - unfair workload assignment, and

Motivation

Learning to defer has limitations

- requires *all* human experts must annotate every training sample
  - impractical (e.g., each sample is annotated by few human experts), and
  - costly, time-consuming, and even infeasible (e.g., radiology),
- most likely selects the *best human expert* all the time
  - unfair workload assignment, and
  - fatigue, burnout $\rightarrow$ misdiagnosis.

Probabilistic L2D - missing annotations

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Each sample must be annotated by **all** human experts.

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Each sample must be annotated by **all** human experts.

$\downarrow$

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Each sample must be annotated by **all** human experts.

$$\downarrow$$

Each sample is annotated by **some** (or none) human experts.

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Each sample must be annotated by **all** human experts.

$\downarrow$

Each sample is annotated by **some** (or none) human experts.

$\rightarrow$ additional latent variables (i.e., missing annotations $\mathbf{t}^{(j)}, \forall j \in \mathcal{D}^{\text{unobs.}}$)

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Each sample must be annotated by **all** human experts.

$$\downarrow$$

Each sample is annotated by **some** (or none) human experts.

$\rightarrow$ additional latent variables (i.e., missing annotations $\mathbf{t}^{(j)}, \forall j \in \mathcal{D}^{\text{unobs.}}$)

Parameter inference is performed through the *variational* EM algorithm:

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Each sample must be annotated by **all** human experts.

$\downarrow$

Each sample is annotated by **some** (or none) human experts.

$\rightarrow$ additional latent variables (i.e., missing annotations $\mathbf{t}^{(j)}, \forall j \in \mathcal{D}^{\text{unobs.}}$)

Parameter inference is performed through the *variational* EM algorithm:

E-step: approximate the posterior $q(\mathbf{z}, \prod_{j \in \mathcal{D}^{\text{unobs.}}} \mathbf{t}^{(j)})$ via variational inference

Probabilistic L2D - missing annotations

Relax the strong assumption in standard L2D

Each sample must be annotated by **all** human experts.

$\downarrow$

Each sample is annotated by **some** (or none) human experts.

$\rightarrow$ additional latent variables (i.e., missing annotations $\mathbf{t}^{(j)}, \forall j \in \mathcal{D}^{\text{unobs.}}$)

Parameter inference is performed through the *variational* EM algorithm:

E-step: approximate the posterior $q(\mathbf{z}, \prod_{j \in \mathcal{D}^{\text{unobs.}}} \mathbf{t}^{(j)})$ via variational inference

M-step: maximise the "completed"-data log-likelihood w.r.t. γ and $\{\theta_m\}_{m=1}^{M+1}$.

Probabilistic L2D - Control workload assignment

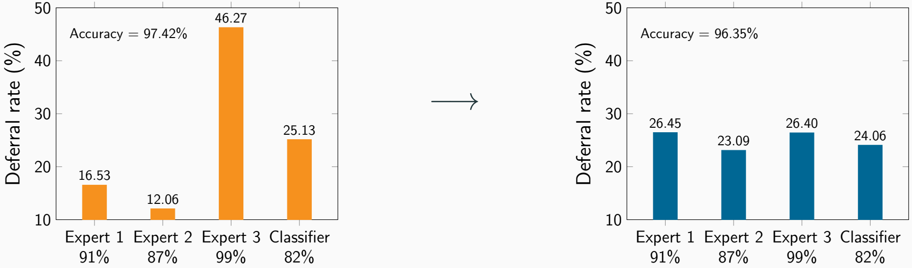

(a) Imbalanced workload

(b) Balanced workload

Figure 2: Workload distribution on Chaoyang dataset at coverage of 0.25.

Probabilistic L2D - Control workload assignment

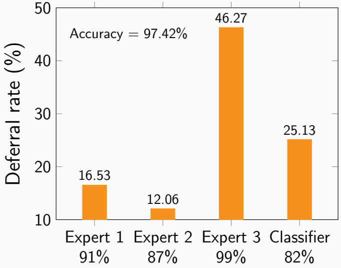

(a) Imbalanced workload

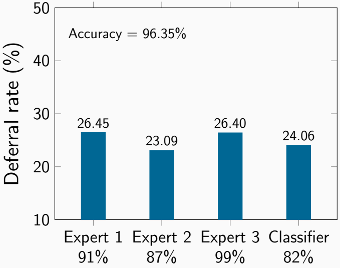

(b) Balanced workload

Figure 2: Workload distribution on Chaoyang dataset at coverage of 0.25.

An additional E-step is introduced to calculate the constrained posterior:

$$\tilde{q}_i^* = \operatorname*{argmin}_{\tilde{q}} \mathrm{KL}\left[\tilde{q}(\mathbf{z}_i) \| q^*(\mathbf{z}_i)\right], \forall i \in \{1, \dots, N\} \quad \text{s.t.: } \varepsilon_\mathrm{l} \preceq \frac{1}{N} \sum_{i=1}^{N} \tilde{q}(\mathbf{z}_i) \preceq \varepsilon_\mathrm{u},$$

where q^* and $\tilde{q}$ denote the unconstrained and constrained posteriors of $\mathbf{z}$.

Evaluation - Coverage - accuracy curve on Chaoyang

Evaluation - Coverage - accuracy curve on Chaoyang

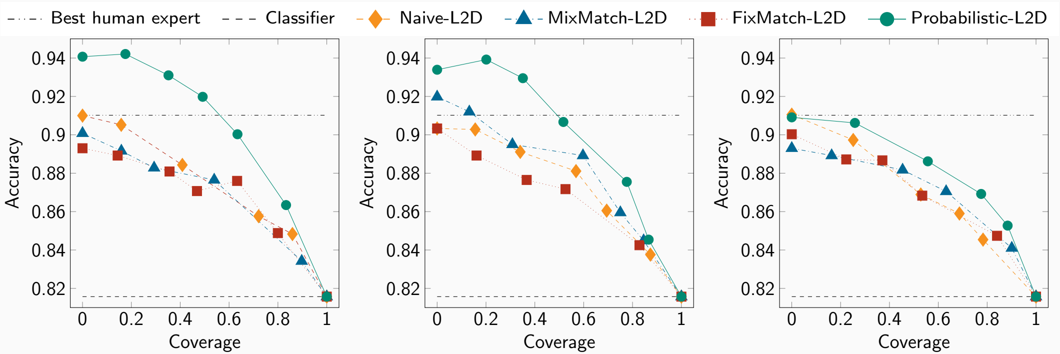

(a) Missing rate of 0.3 **(b)** Missing rate of 0.5 **(c)** Missing rate of 0.7

Figure 3: Comparison of coverage - accuracy curves between different L2D methods on Chaoyang with 2 human experts, each at a different missing rate.

Evaluation - Controllable workload

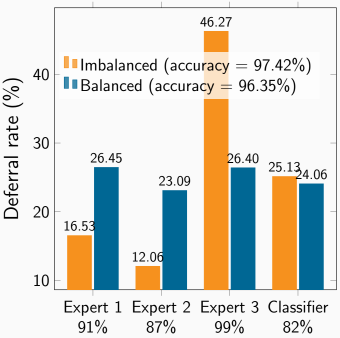

(a) Chaoyang - coverage of 0.25.

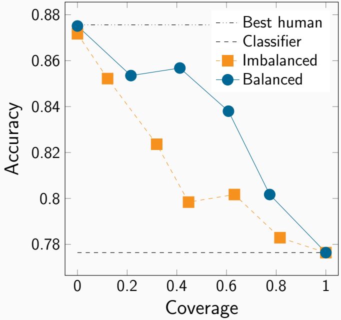

(b) MiceBone - missing rate of 0.3

Figure 4: ((a)) shows comparisons of two different workload constraints on Chaoyang dataset with 50% missing annotations per expert, where in the *imbalanced* setting, $\varepsilon_l = 0$ and $\varepsilon_u = 1$ for each human expert, while in the *balanced* setting, $\varepsilon_l \approx \varepsilon u = (1 - \text{coverage})/M$ for each human expert, and ((b)) coverage - accuracy curve on MiceBone at 30% missing rate.

Conclusion

Conclusion

Propose and develop the probabilistic L2D, which

Conclusion

Propose and develop the probabilistic L2D, which

- addresses the missing annotations, and

Conclusion

Propose and develop the probabilistic L2D, which

- addresses the missing annotations, and
- manages the workload between experts.

Conclusion

Propose and develop the probabilistic L2D, which

- addresses the missing annotations, and
- manages the workload between experts.

Limitations

Conclusion

Propose and develop the probabilistic L2D, which

- addresses the missing annotations, and
- manages the workload between experts.

Limitations

- Scalability w.r.t. the number of human experts

Conclusion

Propose and develop the probabilistic L2D, which

- addresses the missing annotations, and
- manages the workload between experts.

Limitations

- Scalability w.r.t. the number of human experts
- Dynamic expert's performance