Pareto Manifold Learning:

Tackling multiple tasks via ensembles of single-task models

Nikolaos Dimitriadis, Pascal Frossard, François Fleuret

International Conference on Machine Learning

Honolulu - July 23-29, 2023

EPFL

UNIVERSITÉ DE GENÈVE

ICML
International Conference On Machine Learning

Problem formulation

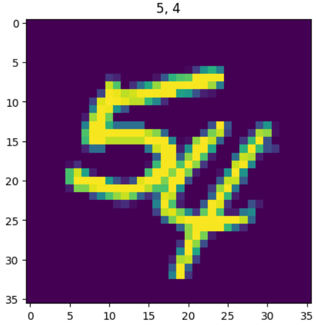

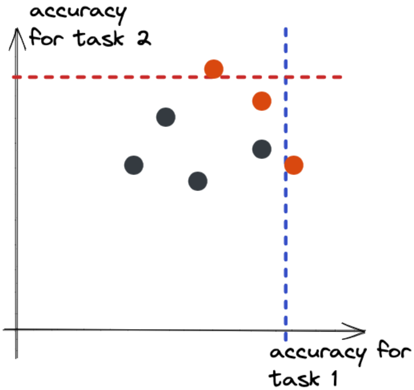

Problem formulation

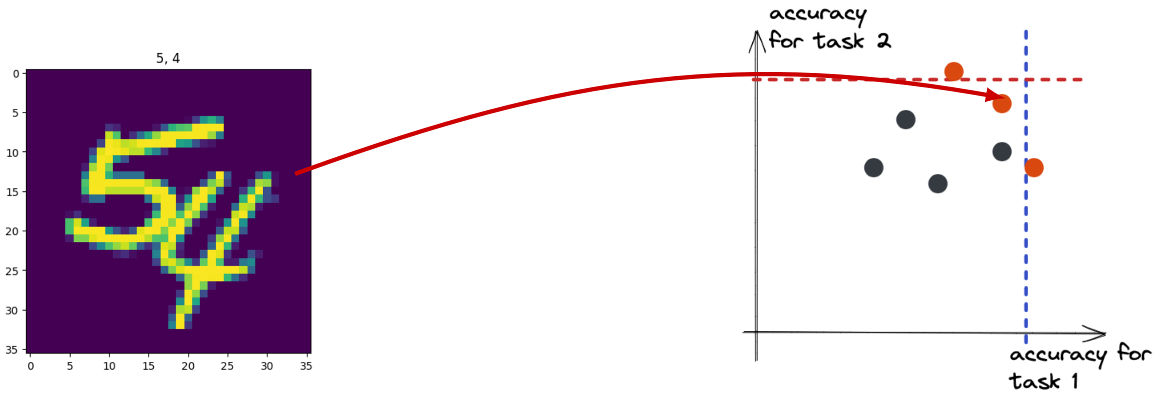

Problem formulation

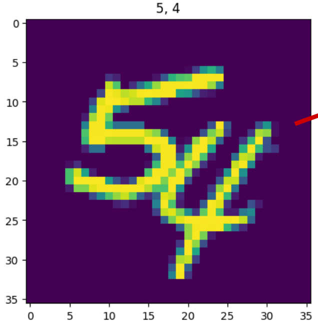

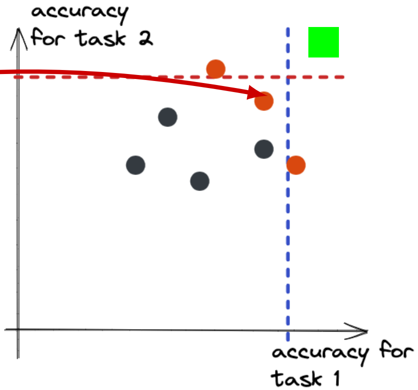

Problem formulation

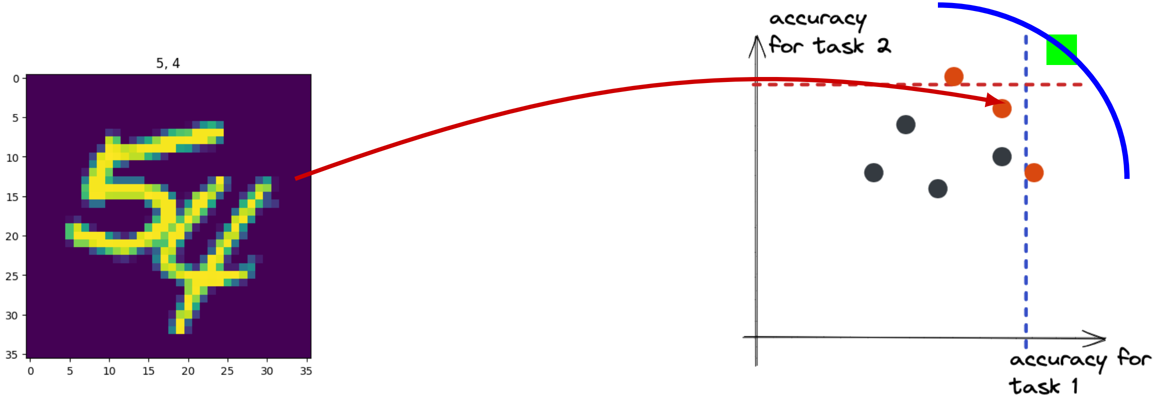

$$\text{ERM objective} = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[\boldsymbol{L}\left(\boldsymbol{y},\boldsymbol{f}\left(\boldsymbol{x};\boldsymbol{\theta}\right)\right)\right] \rightarrow \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[\boldsymbol{w}^\top\boldsymbol{L}\left(\boldsymbol{y},\boldsymbol{f}\left(\boldsymbol{x};\boldsymbol{\theta}\right)\right)\right]$$

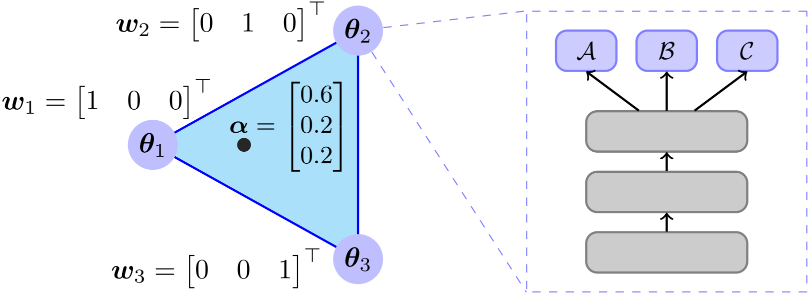

$$\text{ERM objective} = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[\boldsymbol{L}\left(\boldsymbol{y}, \boldsymbol{f}\left(\boldsymbol{x};\boldsymbol{\theta}\right)\right)\right] \to \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[\boldsymbol{w}^\top \boldsymbol{L}\left(\boldsymbol{y}, \boldsymbol{f}\left(\boldsymbol{x};\boldsymbol{\theta}\right)\right)\right]$$

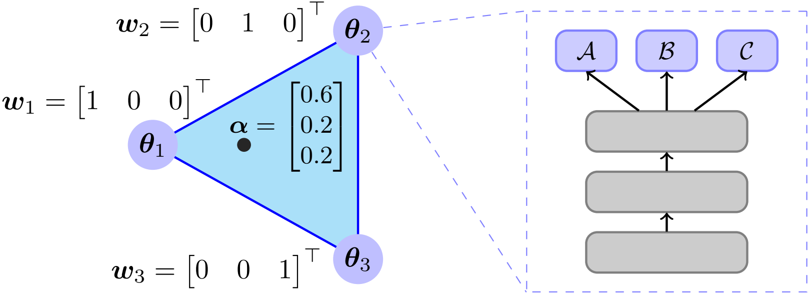

$$\text{objective} = \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[\mathbb{E}_{\boldsymbol{\alpha}\sim\mathrm{P}}\left[\boldsymbol{\alpha}^\top \boldsymbol{L}\left(\boldsymbol{y}, \boldsymbol{f}\left(\boldsymbol{x};\boldsymbol{\alpha}\boldsymbol{\Theta}\right)\right)\right]\right]$$

$$= \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}}\left[\mathbb{E}_{\boldsymbol{\alpha}\sim\mathrm{P}}\left[\sum_{t=1}^{T}\alpha_t\mathcal{L}_t\left(\boldsymbol{y}, \boldsymbol{f}\left(\boldsymbol{x};\sum_{t=1}^{T}\alpha_t\boldsymbol{\theta}_t\right)\right)\right]\right]$$

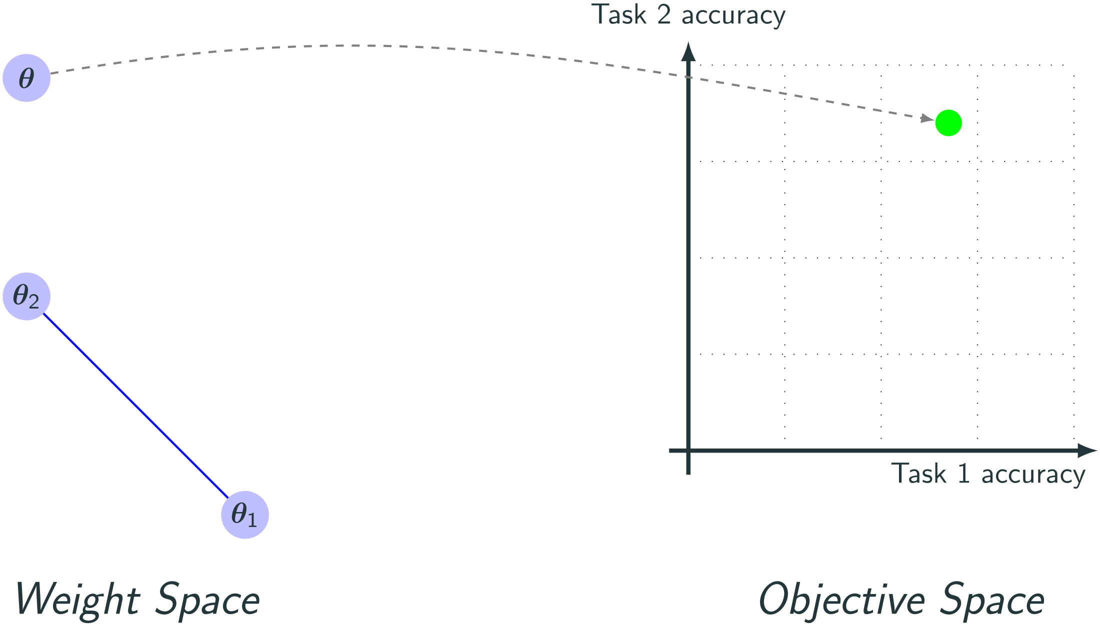

$\boldsymbol{\theta}$
$\boldsymbol{\theta}_2$
$\boldsymbol{\theta}_1$
Task 2 accuracy
Task 1 accuracy
Weight Space
Objective Space

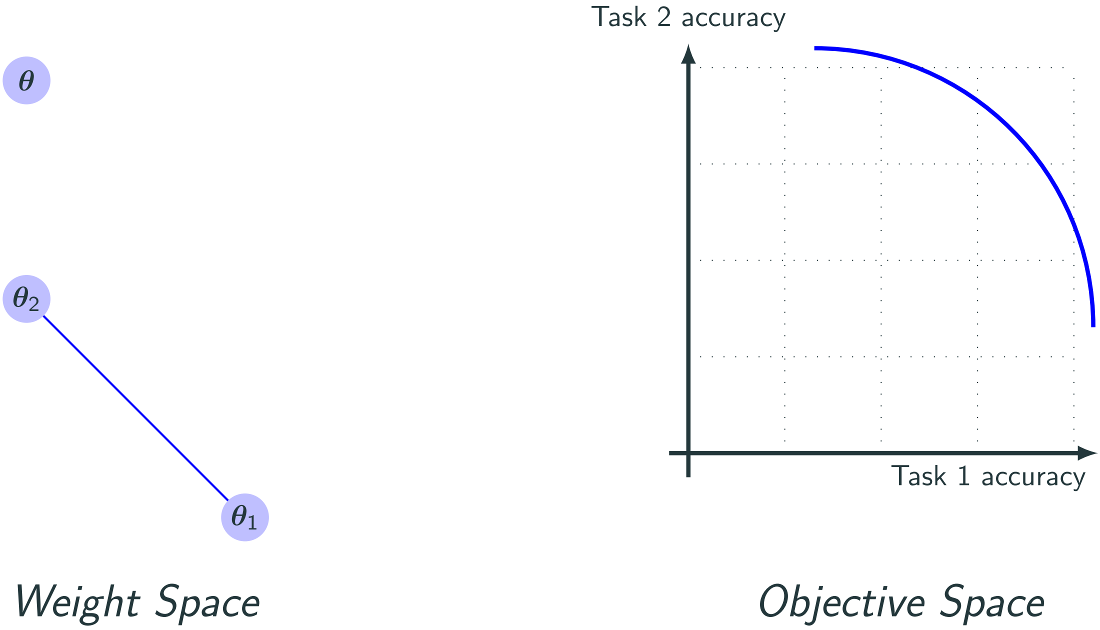

θ
θ_2
θ_1
Weight Space
Task 2 accuracy
Task 1 accuracy
Objective Space

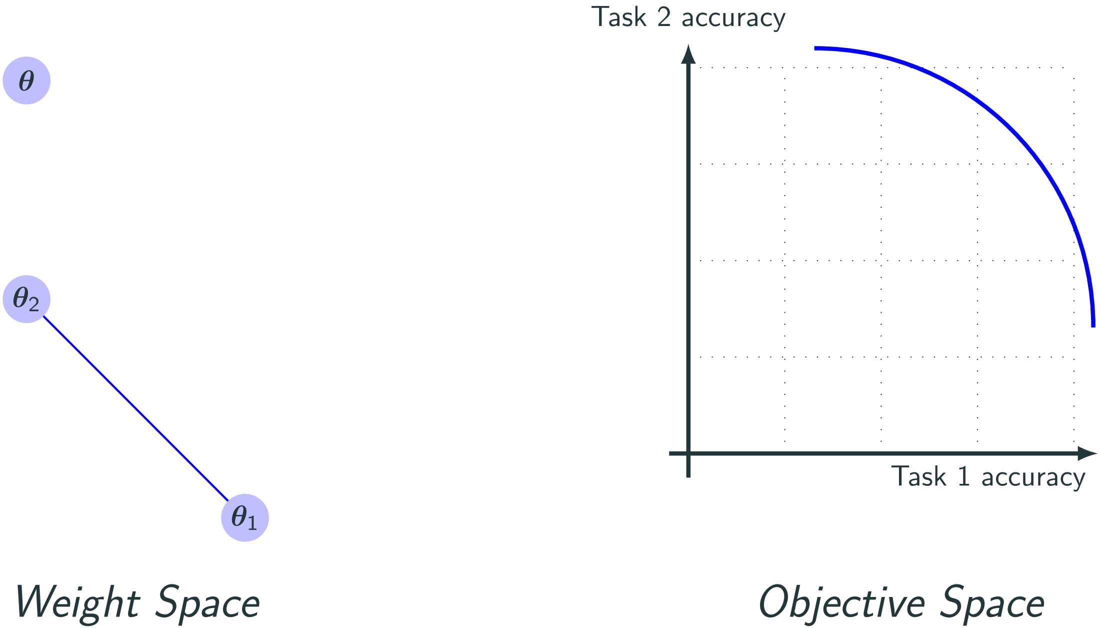

θ
θ_2
θ_1
Weight Space
Task 2 accuracy
Task 1 accuracy
Objective Space

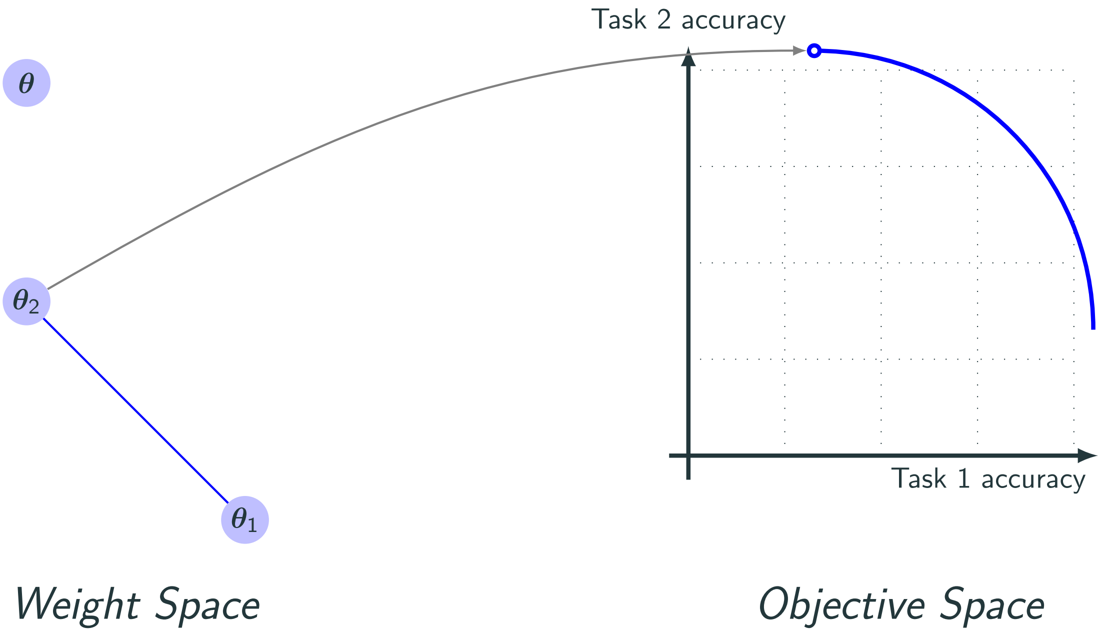

θ
θ_2
θ_1
Task 2 accuracy
Task 1 accuracy
Weight Space
Objective Space

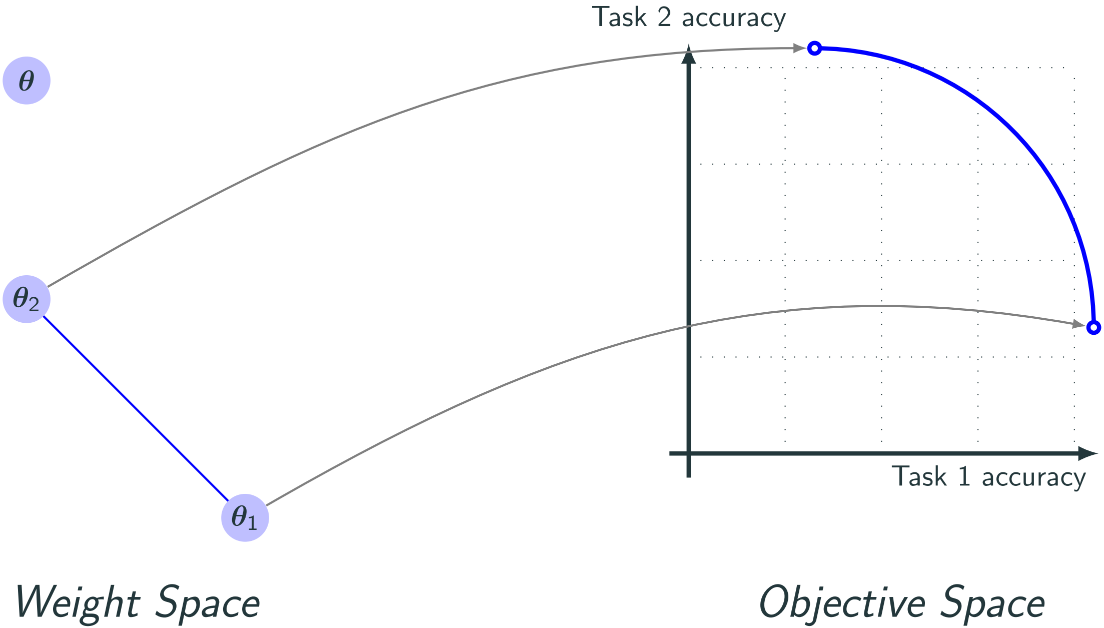
θ
θ_2
θ_1
Task 2 accuracy
Task 1 accuracy
Weight Space
Objective Space

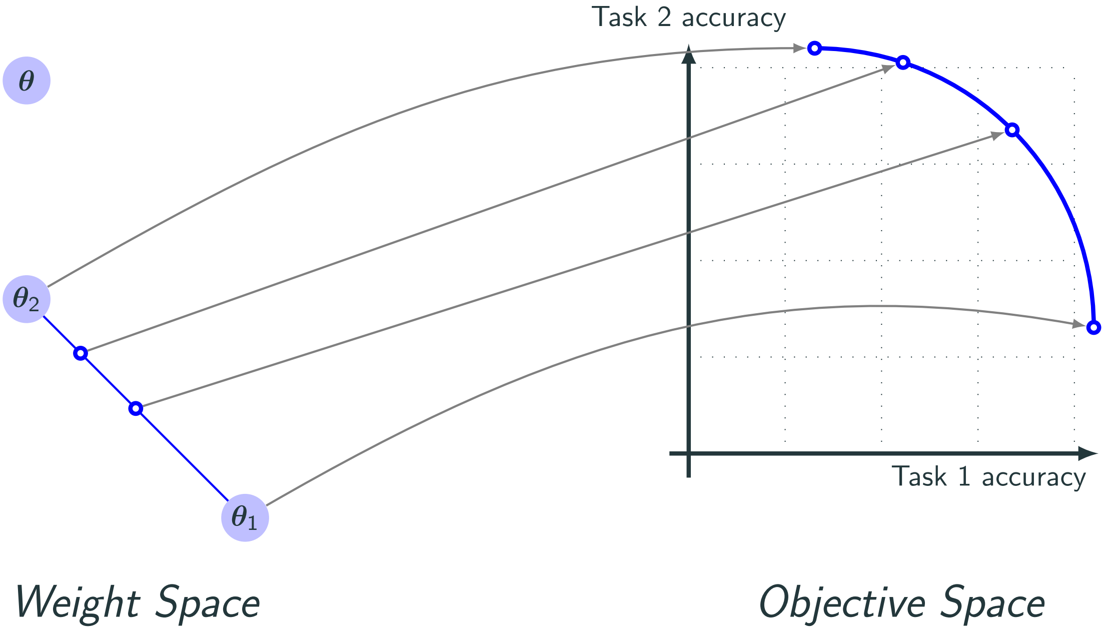

Task 2 accuracy
θ
θ_2
θ_1
Task 1 accuracy
Weight Space
Objective Space

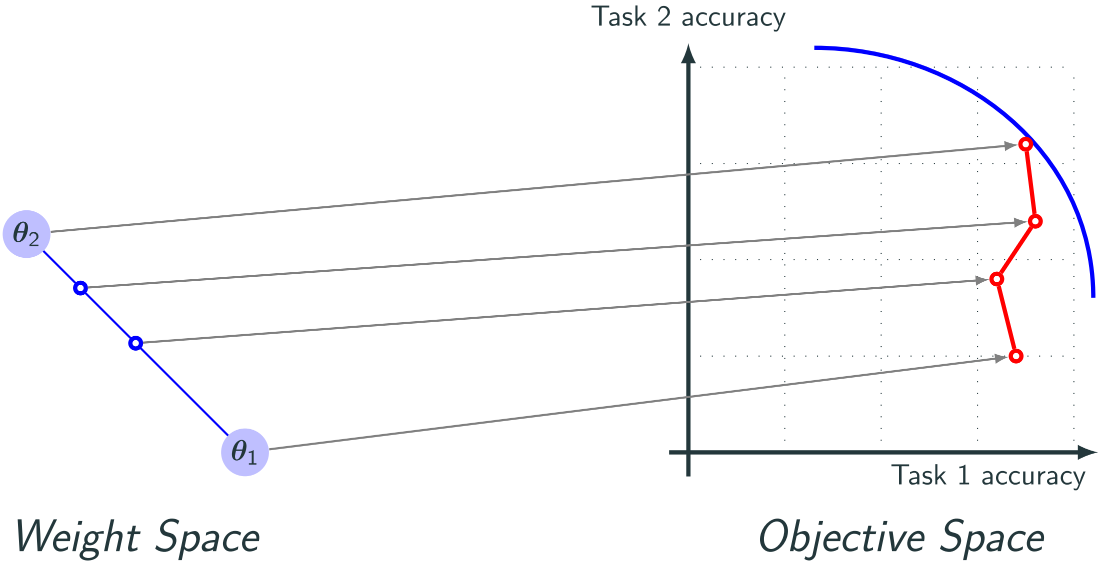

Task 2 accuracy
θ_2
θ_1
Task 1 accuracy
Weight Space
Objective Space

Algorithm Improvements

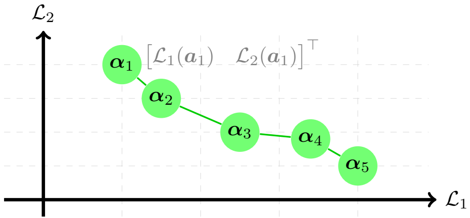

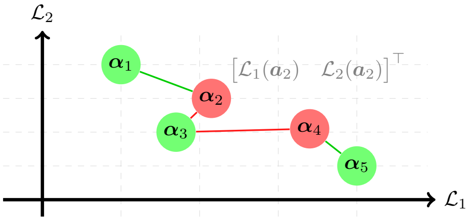

We penalize the violations of the monotonicity constraints for the appropriate task loss.

Algorithm Improvements

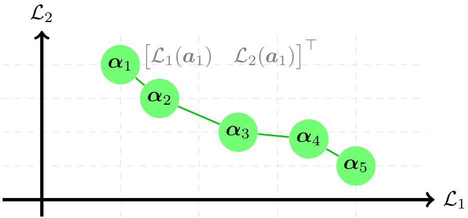

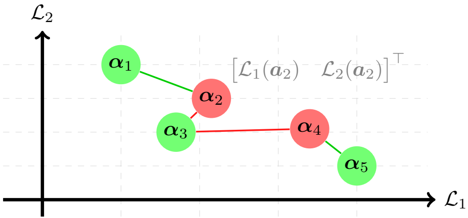

We penalize the violations of the monotonicity constraints for the appropriate task loss.

- Loss and Gradient Balancing Schemes
- Sampling Distribution

Experiments

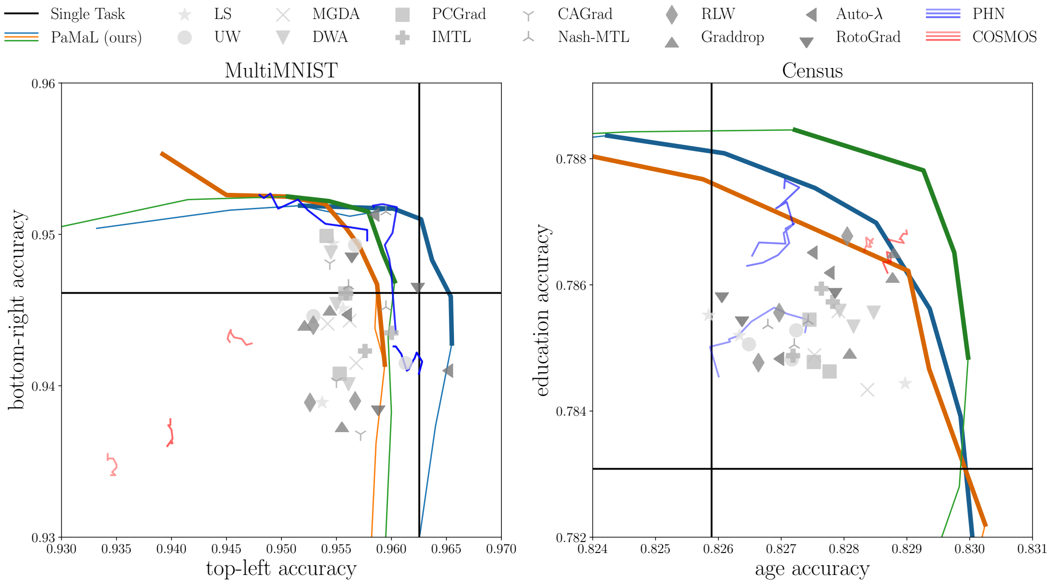

Experiments on `MultiMNIST` and `Census`. Top right is optimal. Three seeds per method.

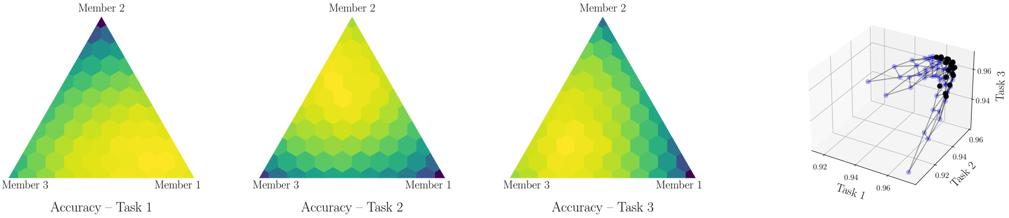

MultiMNIST-3: Accuracy Heatmap and Pareto Front for all tasks.

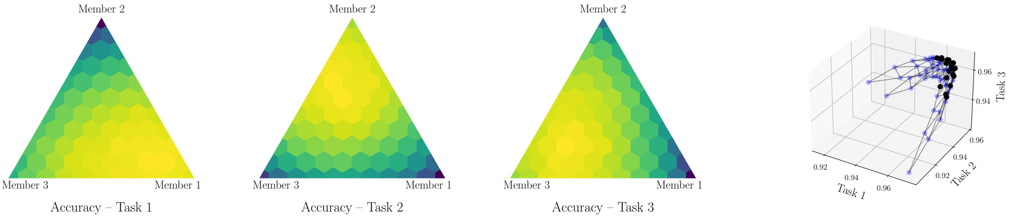

`MultiMNIST-3`: Accuracy Heatmap and Pareto Front for all tasks.

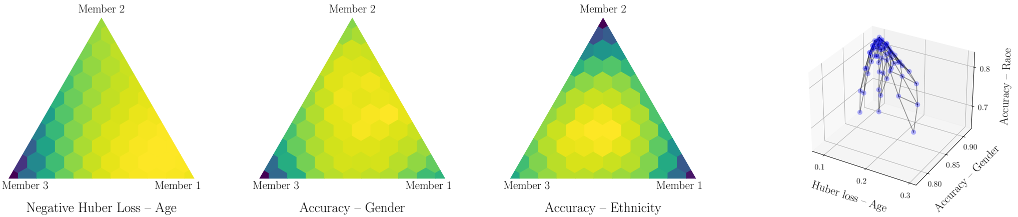

`UTKFace`: Objective Heatmap and Pareto Front for all tasks.

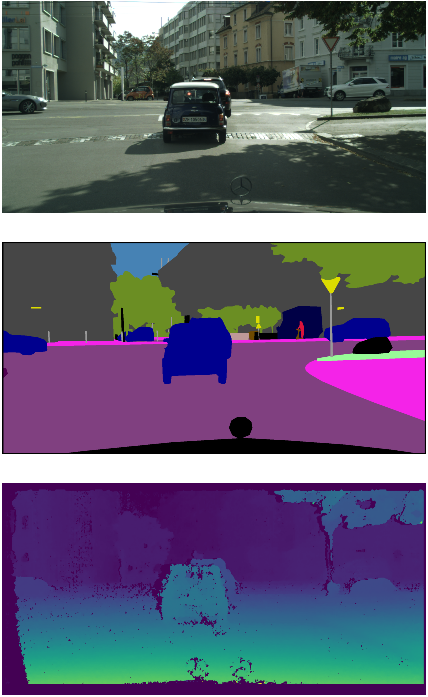

Test performance on *CityScapes*. 3 random seeds per method.

	Segmentation		Depth	
	mIoU ↑	Pix Acc ↑	Abs Err ↓	Rel Err ↓
STL	70.96	92.12	**0.0141**	38.644
LS	70.12	91.90	0.0192	124.061
UW	70.20	91.93	0.0189	125.943
MGDA	66.45	90.79	**0.0141**	53.138
DWA	70.10	91.89	0.0192	127.659
PCGrad	70.02	91.84	0.0188	126.255
IMTL	70.77	92.12	0.0151	74.230
Graddrop	70.07	91.93	0.0189	127.146
CAGrad	69.23	91.61	0.0168	110.139
RLW	68.79	91.52	0.0213	126.942
Nash-MTL	**71.13**	**92.23**	0.0157	78.499
RotoGrad	69.92	91.85	0.0193	127.281
Auto-λ	70.47	92.01	0.0177	116.959
COSMOS	69.78	91.79	0.0539	136.614
PaMaL(ours)	70.35	91.99	**0.0141**	54.520

Pareto Manifold Learning:

Tackling multiple tasks via ensembles of single-task models

Nikolaos Dimitriadis, Pascal Frossard, François Fleuret

International Conference on Machine Learning

Honolulu - July 23-29, 2023

EPFL

UNIVERSITÉ DE GENÈVE

ICML International Conference On Machine Learning