

NEURAL INFORMATION PROCESSING SYSTEMS

UNICAMP

UNIVERSITÄT DES SAARLANDES

erc European Research Council
Established by the European Commission

Learning Pareto Stationary Fronts via Single-Pass Backpropagation

Elina Rojin Celik



Marcos Medeiros Raimundo



Isabel Valera



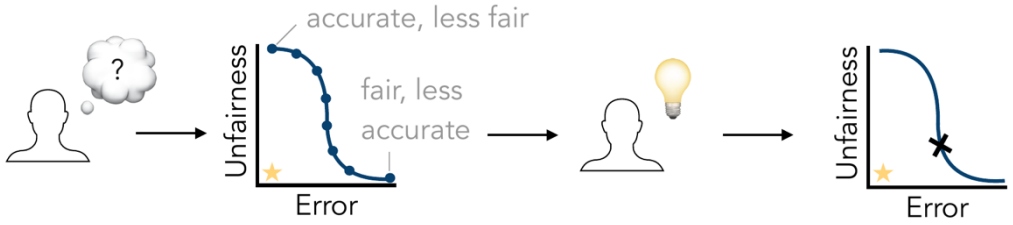
Balancing competing objectives

UNIVERSITÄT DES SAARLANDES

STRONG CONFLICT

Fairness vs. accuracy

Improving one objective systematically degrades the other.

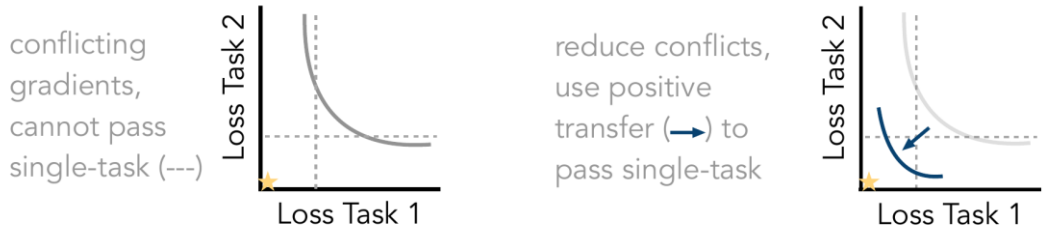


Goal: recover diverse solutions to make an informed choice

WEAK CONFLICT

Multi-task learning

Positive transfer is possible, but gradient conflicts lead to suboptimal solutions.



Goal: use positive transfer to move closer to utopia

What we want from a method — so far, none satisfies all three

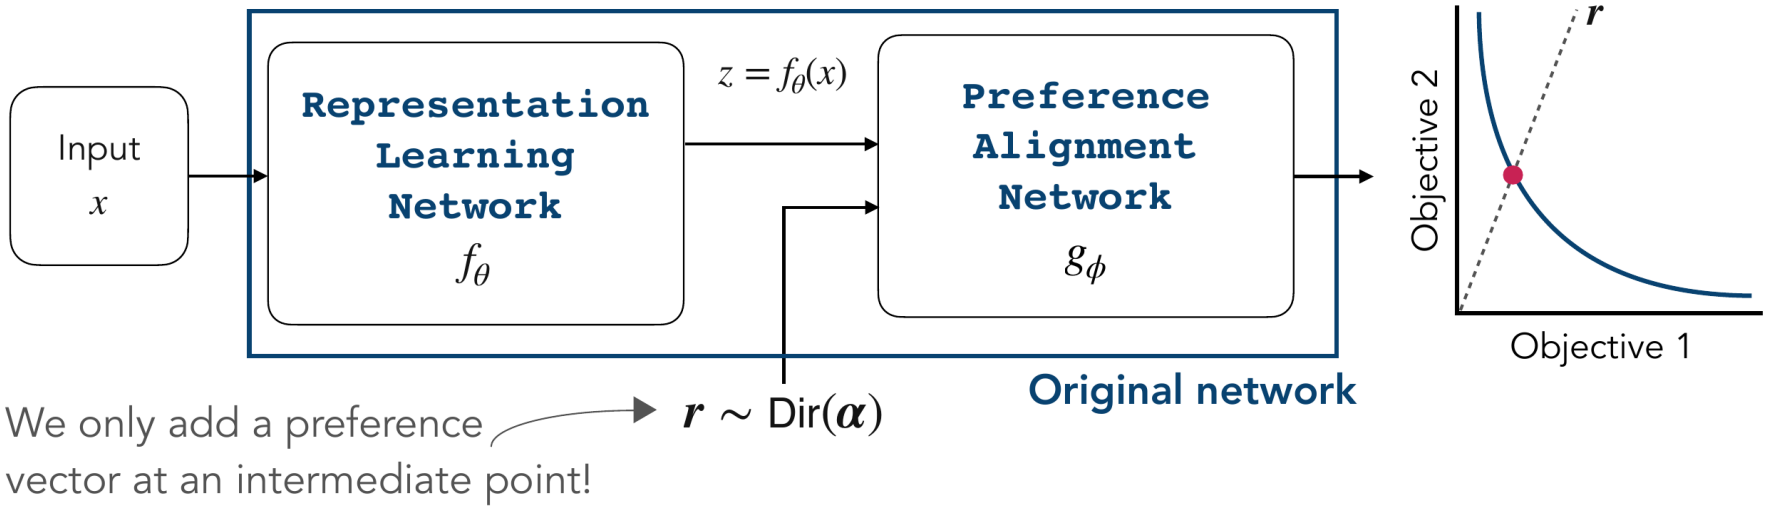
1 **Efficiency**
No extra cost over standard single-objective training

2 **Optimality**
Solutions are Pareto stationary

3 **Diversity**
Captures the full spectrum of trade-offs

MOSEL: one network for every trade-off

MOSEL (Multi-Objective Stackelberg Efficient Learning) recovers full **Pareto stationary fronts** at the **cost of standard single-objective training**.



Bilevel formulation

- **Leader:** preference alignment network (upper level)
- **Follower:** representation learning network (lower level)

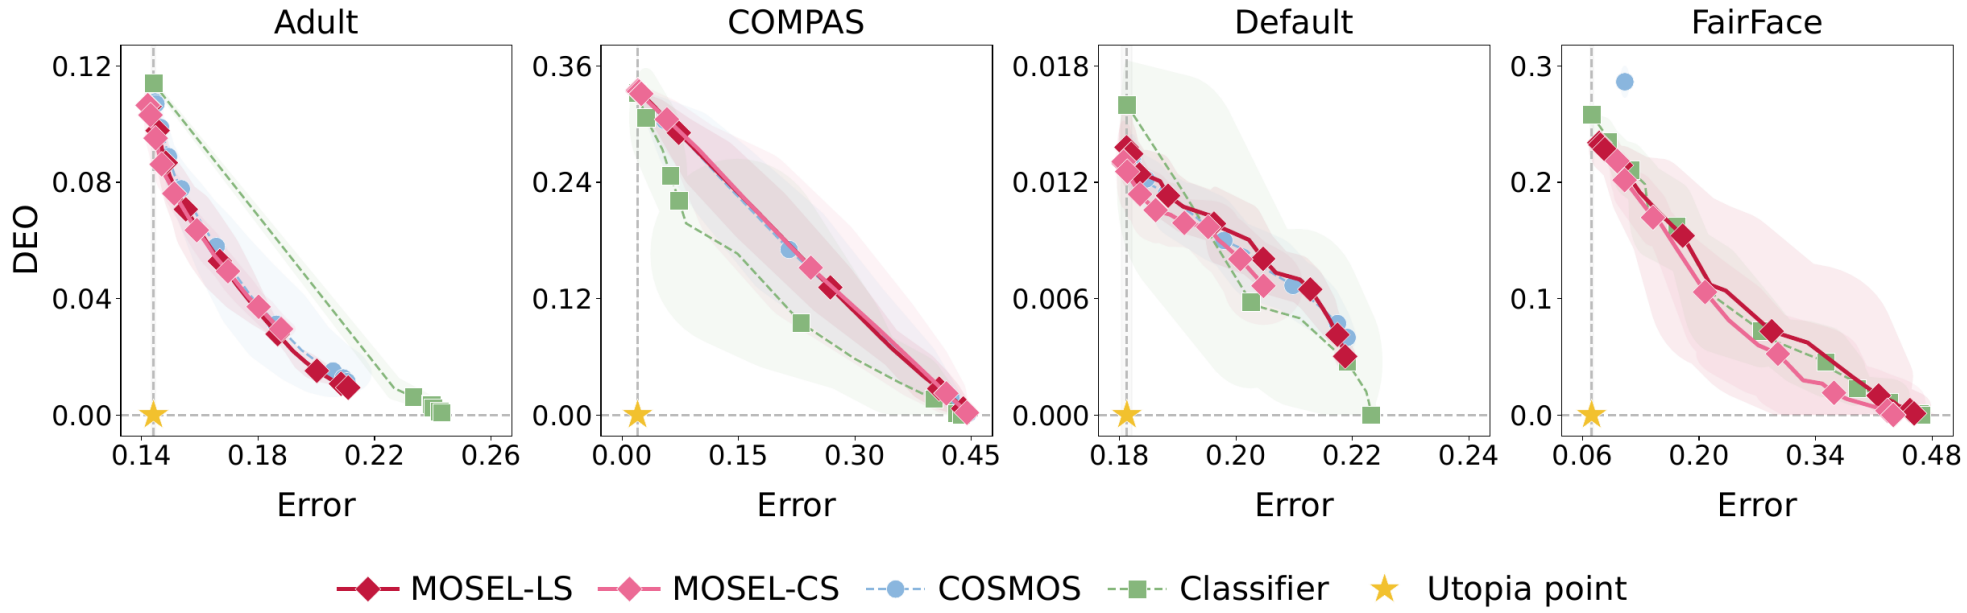
One forward-backward pass

- Sample one preference vector per batch
- Update both networks; using a smaller learning rate for preference alignment mimics the hierarchy

Theory

- Linear & Chebyshev scalarization: solutions are Pareto stationary
- Sampling acts as implicit bias toward shared features

Strong conflict: fairness vs. accuracy

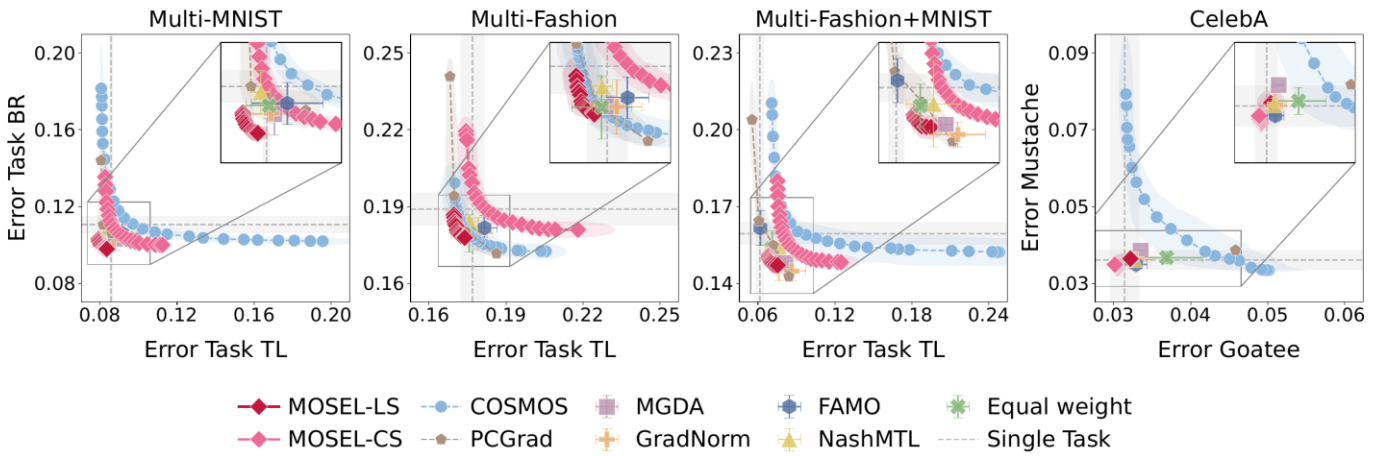


Hypervolume ↑	Adult	COMPAS	Default	FairFace
MOSEL-LS	**.849**	.568	**.922**	**.701**
MOSEL-CS	.776	.560	.897	**.734**
COSMOS	.830	.567	.905	.176
Post-hoc	.618	**.645**	**.925**	**.708**

✓ Diverse Pareto fronts

MOSEL traces the **full trade-off** in strong-conflict settings and is the only learned front that stays competitive on all four datasets.

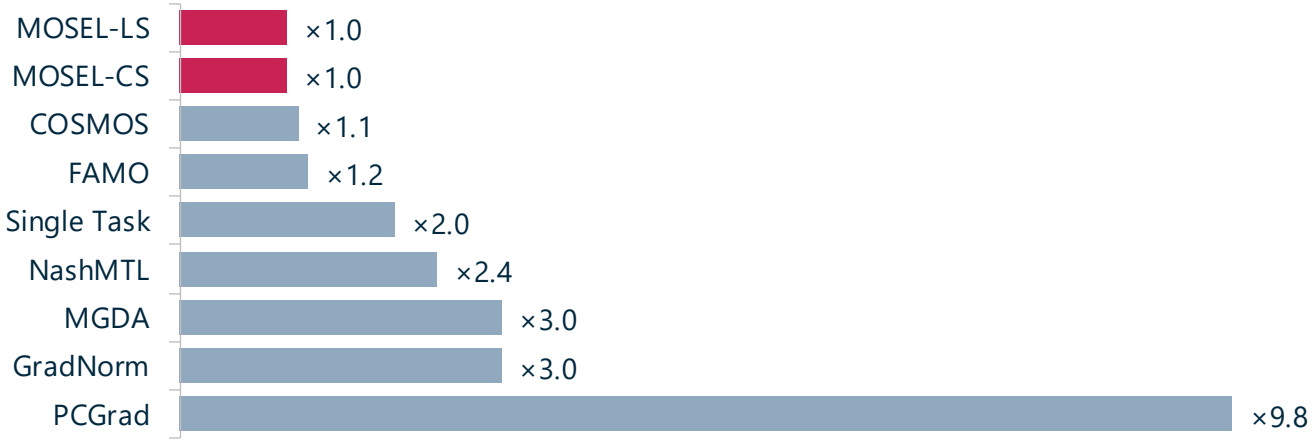
Weak conflict: multi-task learning



✓ Fewer gradient conflicts

Solutions sit **closer to utopia** than state-of-the-art MTL methods, at the **cost of standard single-objective training.**

Training time on CelebA (relative to equal weighting)



Scalability on CelebA

- Up to **10 objectives** at roughly constant training time
- Correlated tasks: up to **+2.2%** over single-task, growing with more objectives
- Uncorrelated tasks: most objectives still improve

Only 672 extra parameters on a 17.5M-parameter model

Multi-Objective Stackelberg Efficient Learning

1 **Efficiency**
No extra cost over standard single-objective training

2 **Optimality**
Solutions are Pareto stationary

3 **Diversity**
Captures the full spectrum of trade-offs







erc European Research Council
Established by the European Commission

UNICAMP

UNIVERSITÄT DES SAARLANDES