Distribution-First Framework for Learning Risk-Sensitive Individualized Treatment Rules

Zhen Fang[1]*, Guanting Chen[1], Yufeng Liu[2]

- [1]Department of Statistics and Operations Research, University of North Carolina at Chapel Hill
- [2]Department of Statistics, University of Michigan





THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL



Individualized Decision Making

- Example: personalized medicine
  - Tailor each patient's therapy to their covariates (age, biomarkers, history) to improve health outcomes
- ★ Individualized Treatment Rule (ITR) d: covariates $x \mapsto$ treatment a
- Classical goal: maximize the expected outcome [Kosorok and Laber, 2019]

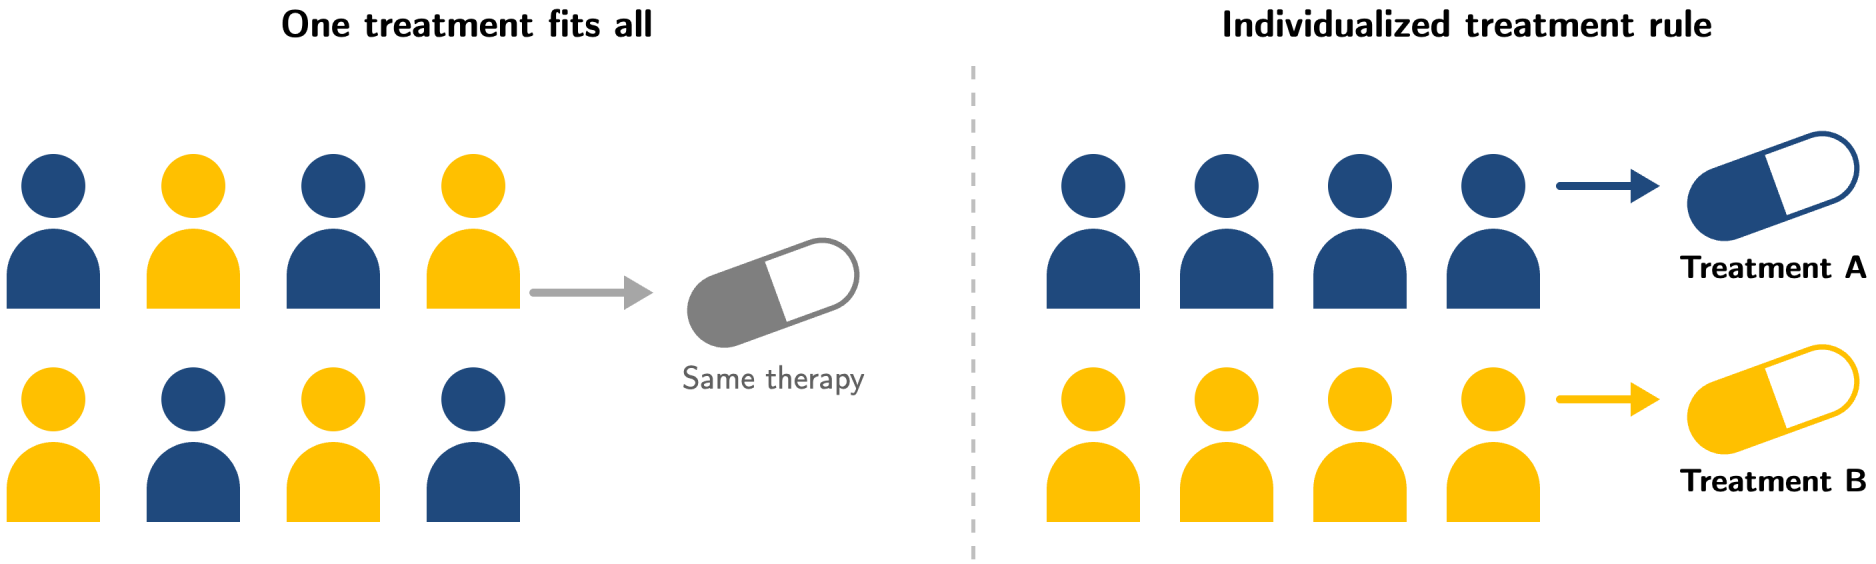


Figure 1: Transition from "one size fits all" to individualized treatment.

Mean not enough, distribution matters

★ A mean-based ITR **cannot** tell A and B apart!

- Two treatments with the **same average** outcome
- Treatment B: most patients do slightly better, but a few do **very badly** (lower tail)

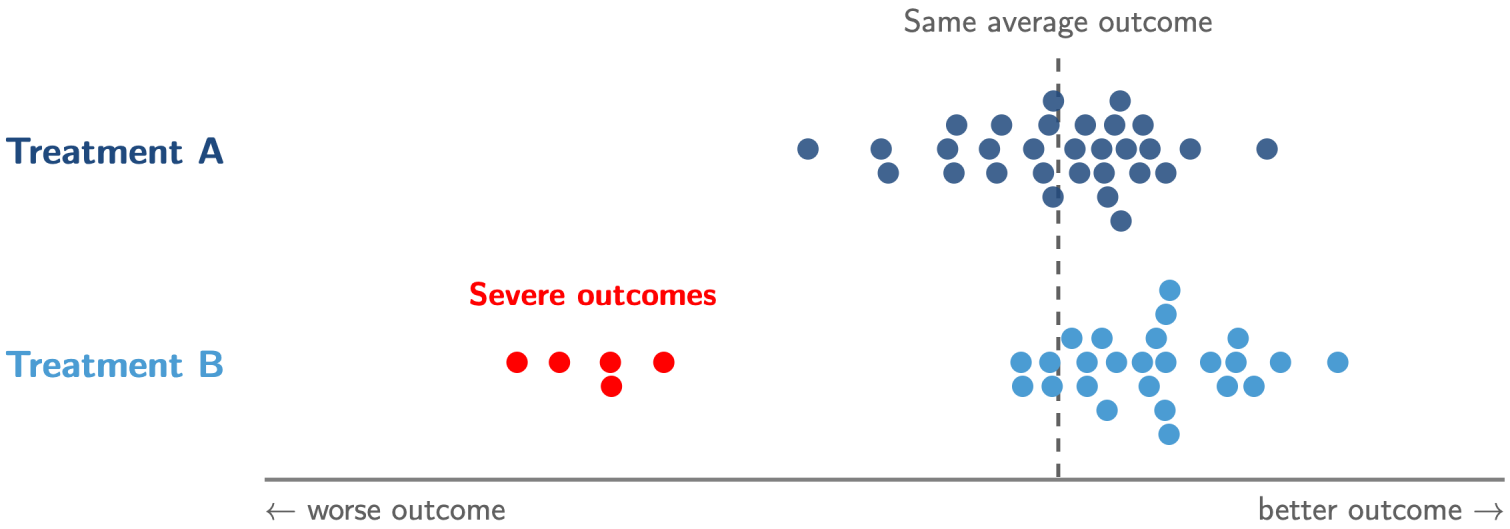


Figure 2: Illustration — the conditional mean hides the lower tail (each dot = one patient)

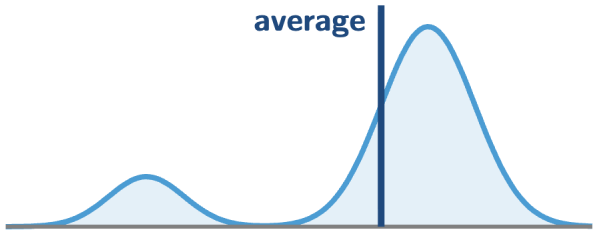
THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Different Patients, Different Criteria

- The "best" treatment depends on how a patient weighs risk

Risk-neutral

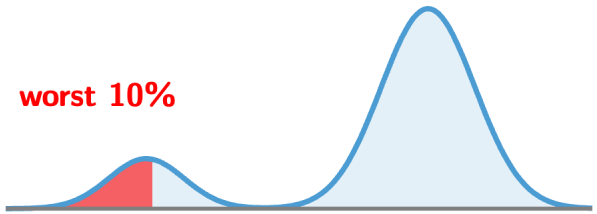
→ Mean



Maximize the **average** outcome

Risk-averse

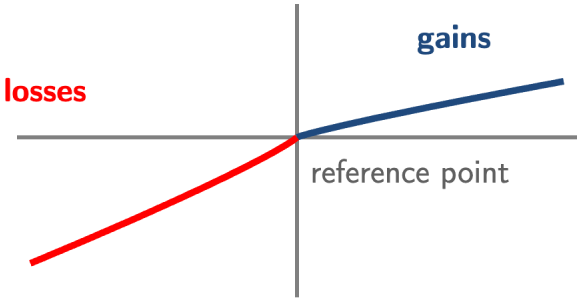
→ CVaR



Maximize the average of the worst 10% of outcomes

Loss-averse

→ CPT



Losses loom larger than gains [Tversky and Kahneman, 1992]

Figure 3: Three patients, three treatment criteria, one outcome distribution.

★ Same distribution, different criteria ⇒ possibly **different** optimal treatments

Motivations and Challenges

1. Criterion-specific learning [Wang et al., 2018; Qi et al., 2023]:
   - Each criterion needs its **own** learner and tuning
   - New patient preference ⇒ **retrain** from scratch

   ★ How to compare **many** criteria within one pipeline?

2. Hard-to-optimize objectives:
   - CPT weights cumulative probabilities **nonlinearly**
   - No simple pointwise loss to minimize directly

   ★ How to learn ITRs for **complex** risk-sensitive criteria?

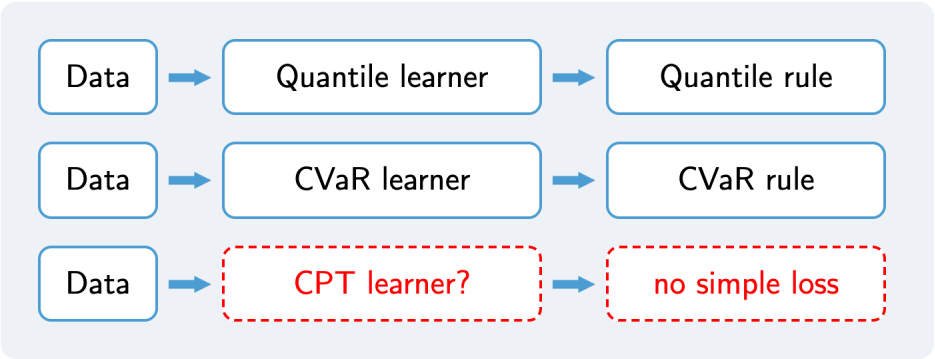


Figure 4: One pipeline per criterion.

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Distribution-First Framework

- ★ Idea: every criterion is a functional of the conditional outcome distribution ⇒ estimate the distribution first

Stage 2: $\widehat{Q}_a(x) = T\{\widehat{F}_a(\cdot \mid x)\}$ Stage 3: $\widehat{\pi}(x) \in \arg\max_a \widehat{Q}_a(x)$

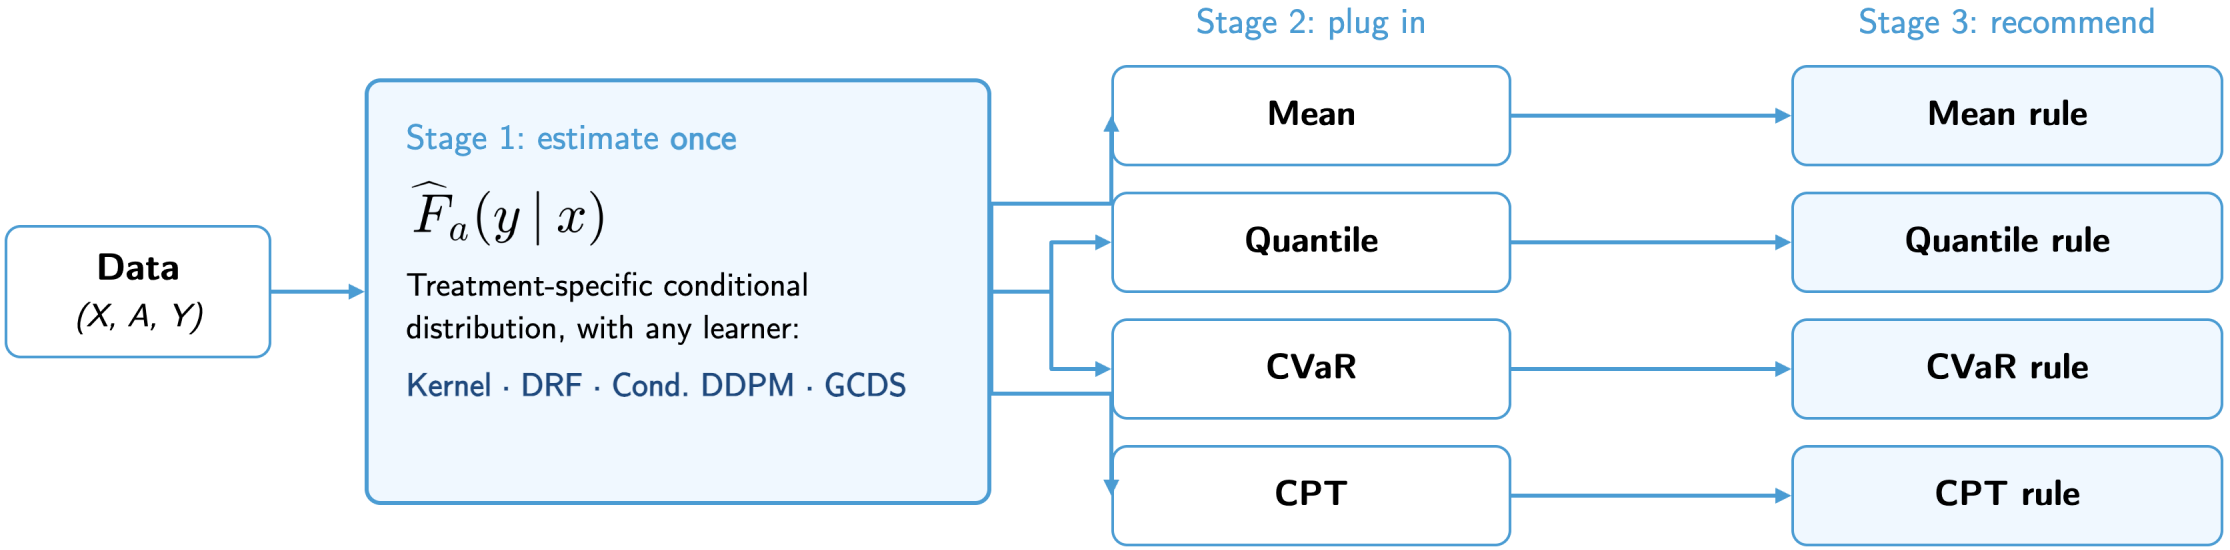


Figure 5: One Stage-1 fit [Ćevid et al., 2022] serves every criterion — no retraining.

- ★ Theory: better distribution estimation ⇒ better decisions (regret guarantee, covering CPT)

41–62% Fewer Wrong Treatment Decisions in Challenging Settings

- Decision error = share of patients given a suboptimal treatment (lower is better)

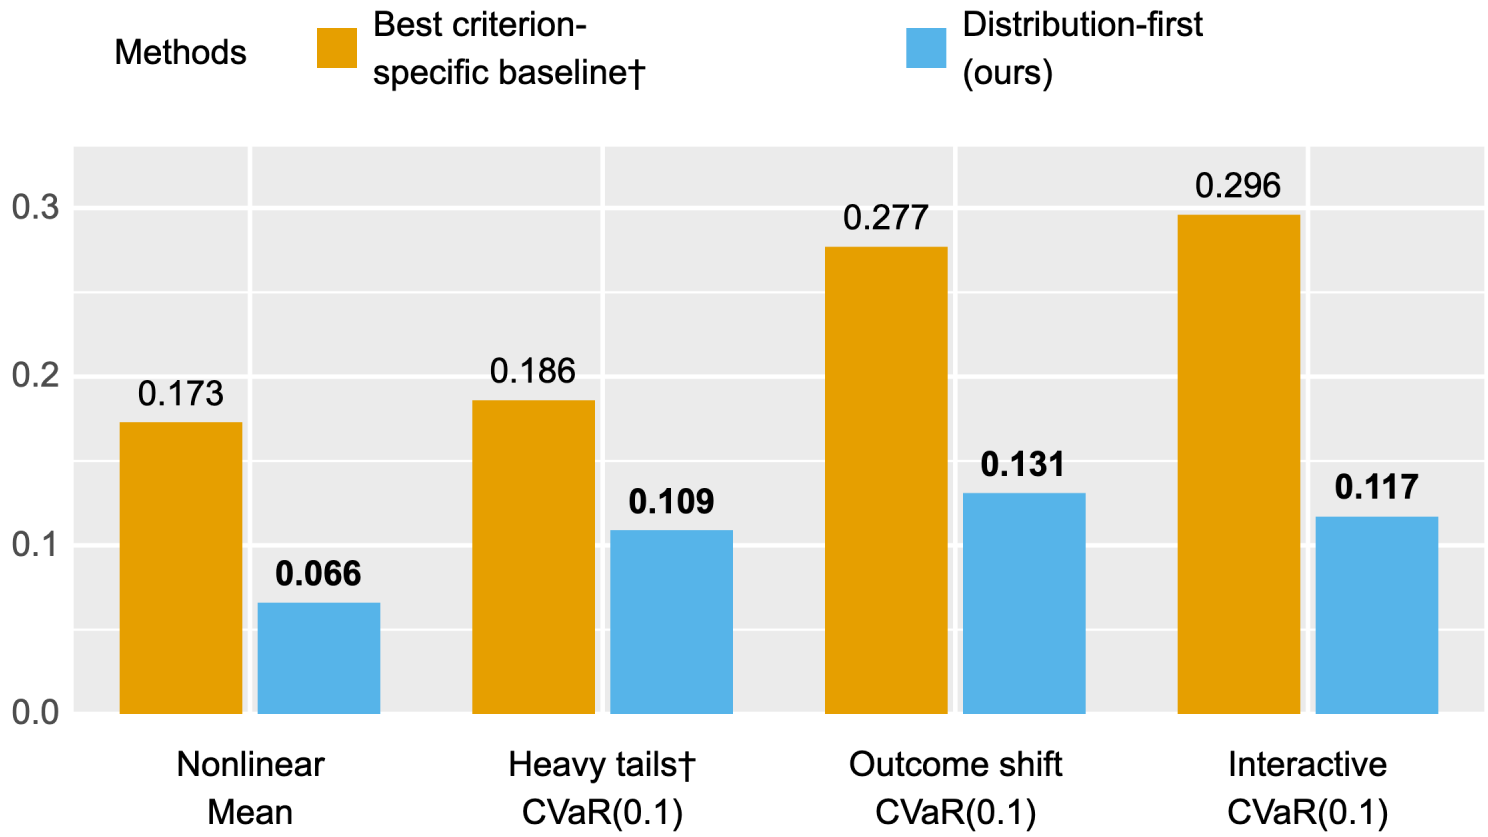


Figure 6: Decision error, mean over 100 replications.

† Excludes linear CQR, which is correctly specified under heavy tails (0.009). Linear Gaussian setting: our best estimator matches the best baseline (0.032 vs 0.034).

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

ACTG175: One Fit, Different Criteria, Different Treatment Decisions

- 1,046 HIV patients [Hammer et al., 1996]: ZDV + ddI vs ZDV + zalcitabine; outcome = change in CD4 count

Method	Target	Mean	CVaR 0.25	CPT	Runtime (s)
ℓ1-PLS	Mean	51.8	−107.9	−3.49	0.07
ℓ2-CVaR	CVaR 0.25	52.9	−108.0	−2.95	193.6†
CQR	Quantile 0.25	48.3	−110.0	−6.12	0.005
DRF (ours)	Mean	54.0*	−108.8	−3.06	4.0
DRF (ours)	CVaR 0.25	45.0	−105.4	−2.56	(shared)
DRF (ours)	CPT	53.6	−104.6*	−1.25*	(shared)

Table 1: Held-out utility over 200 splits (higher is better); * best. Runtime: median seconds per split.

One DRF fit (4.0 s) serves all three rules. † Includes 5-fold CV over 11 tuning values; 1.8 s once tuned.

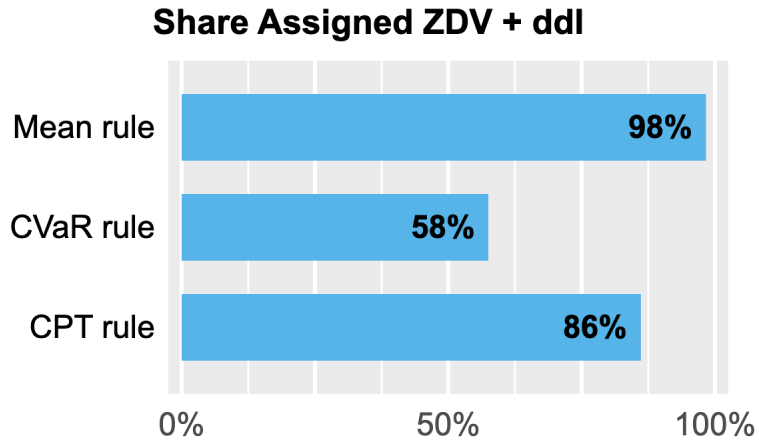


Figure 7: Same 4-second fit, different recommendations.

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Summary

In summary, we

- Proposed a distribution-first framework: estimate the conditional distribution once, plug in any criterion
- Linked Stage-1 distribution error to regret, including non-Lipschitz CPT
- Showed **41–62% fewer** wrong decisions in simulations, and criterion-dependent rules on ACTG175
- Code: github.com/zhenf530/Distributio_Risk_ITR

☺ **Thanks for your listening!**

☼ **Welcome to join our poster session:**

Poster Session 3: 10:00 AEDT, December 9th (Wednesday), Hall 1–4, Sydney

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

References

Ćevid, D., Michel, L., Näf, J., Bühlmann, P., and Meinshausen, N. (2022).
Distributional random forests: Heterogeneity adjustment and multivariate distributional regression.
Journal of Machine Learning Research, 23(333):1–79.

Hammer, S. M., Katzenstein, D. A., Hughes, M. D., et al. (1996).
A trial comparing nucleoside monotherapy with combination therapy in HIV-infected adults with CD4 cell counts from 200 to 500 per cubic millimeter.
New England Journal of Medicine, 335(15):1081–1090.

Kosorok, M. R. and Laber, E. B. (2019).
Precision medicine.
Annual Review of Statistics and Its Application, 6:263–286.

Qi, Z., Pang, J.-S., and Liu, Y. (2023).
On robustness of individualized decision rules.
Journal of the American Statistical Association, 118(543):2143–2157.

Tversky, A. and Kahneman, D. (1992).
Advances in prospect theory: Cumulative representation of uncertainty.
Journal of Risk and Uncertainty, 5(4):297–323.

Wang, L., Zhou, Y., Song, R., and Sherwood, B. (2018).
Quantile-optimal treatment regimes.
Journal of the American Statistical Association, 113(523):1243–1254.



THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL