SPREAD: Sampling-based Pareto Front Refinement via Efficient Adaptive Diffusion

Sedjro Salomon Hotegni*, Sebastian Peitz

Project

One-sentence takeaway

SPREAD turns conditional diffusion into a multi-objective optimizer by combining shifted conditioning, MGD-inspired adaptive guidance, and diversity-promoting repulsion, yielding strong Pareto-front quality in online, offline, and Bayesian settings.

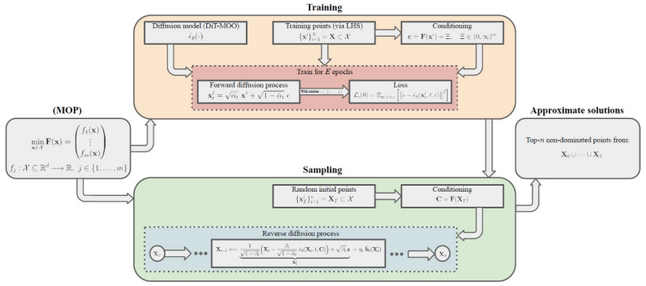

Challenges in Multi-Objective Optimization (MOO)

- Many real-world problems require **optimizing multiple conflicting objectives at once**.
- The goal is not a single optimum, but a Pareto front of trade-off solutions.
- Classical methods can struggle in **high-dimensional**, **large-scale**, and **resource-constrained** settings.

Why diffusion for MOO?

- **Iterative refinement:** diffusion models improve samples step by step, which naturally matches moving candidates toward the Pareto front.
- **Conditional generation:** they can learn a conditional process over decision variables to steer sampling toward promising trade-offs.
- **Optimization-aware sampling:** reverse diffusion can be guided to improve objective values while preserving diversity.

Main contributions

- A **diffusion-based generative framework** for MOO.
- A **novel conditioning strategy** with a theoretical improvement guarantee.
- An **adaptive guidance mechanism** inspired by multiple gradient descent (MGD).
- A **diversity-promoting repulsion** term for better Pareto-front coverage.
- Strong results in **online**, **offline**, and **Bayesian** MOO settings.

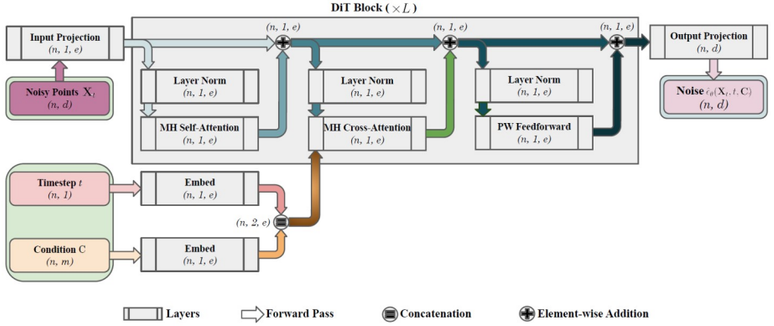

Method overview

- **Shifted conditioning:** For decision vectors $\mathbf{x}^i$, the model is trained with *shifted labels* $\mathbf{c}^i = \mathbf{F}(\mathbf{x}^i) + \Xi, \quad \Xi \in (\mathbf{0}, +\infty)^m$.
- **Guidance objective:** At each reverse step, SPREAD selects directions that balance alignment with *MGD descent*, and dispersion in objective space via a *Gaussian RBF repulsion* term.
- **Final solution set:** Across reverse steps, SPREAD maintains the top non-dominated candidates using *crowding distance* to preserve diversity. The final output is an approximation of the Pareto set gathered over the diffusion trajectory.

Experiments and Results

Online

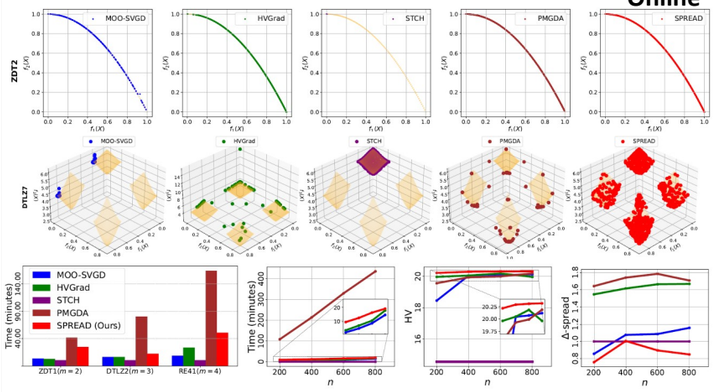

Offline

HV (↑)	$m=2$			$m=3$							
Method	ZDT1	ZDT2	ZDT3	ZDT4	ZDT6	DTLZ1	DTLZ2	DTLZ3	DTLZ4	DTLZ5	DTLZ6	DTLZ7
Ȳ(best)	4.17	4.67	5.15	5.45	4.61	10.60	9.91	10.00	10.76	9.35	8.88	8.56
MM	4.81 ± 0.02	5.57 ± 0.07	5.48 ± 0.21	5.03 ± 0.19	4.78 ± 0.01	10.64 ± 0.01	9.03 ± 0.80	10.58 ± 0.03	7.66 ± 1.30	7.65 ± 1.39	9.58 ± 0.31	10.61 ± 0.16
MM-COM	4.52 ± 0.02	4.99 ± 0.12	5.49 ± 0.07	5.10 ± 0.08	4.41 ± 0.21	10.64 ± 0.01	8.99 ± 0.97	10.27 ± 0.37	9.72 ± 0.39	9.44 ± 0.41	9.37 ± 0.35	10.09 ± 0.36
MM-IOM	4.68 ± 0.12	5.45 ± 0.11	5.61 ± 0.06	4.99 ± 0.21	4.75 ± 0.01	10.64 ± 0.01	10.10 ± 0.27	10.24 ± 0.13	10.03 ± 0.53	9.77 ± 0.18	9.30 ± 0.01	10.60 ± 0.05
MM-ICT	4.82 ± 0.01	5.58 ± 0.01	5.59 ± 0.06	4.63 ± 0.43	4.75 ± 0.01	10.64 ± 0.01	8.68 ± 0.88	10.25 ± 0.42	10.33 ± 0.24	9.25 ± 0.28	9.10 ± 1.16	10.29 ± 0.05
MM-RoMA	4.84 ± 0.01	5.43 ± 0.35	**5.89 ± 0.04**	4.13 ± 0.11	1.71 ± 0.10	10.64 ± 0.01	10.04 ± 0.05	10.61 ± 0.03	9.25 ± 0.11	8.71 ± 0.47	9.84 ± 0.25	10.53 ± 0.04
MM-TriMentoring	4.64 ± 0.10	5.22 ± 0.11	5.16 ± 0.04	**5.12 ± 0.12**	2.61 ± 0.01	10.64 ± 0.01	9.39 ± 0.35	10.48 ± 0.12	10.21 ± 0.06	7.69 ± 1.03	9.00 ± 0.48	10.12 ± 0.09
MH	4.80 ± 0.03	5.57 ± 0.07	5.58 ± 0.20	4.59 ± 0.26	4.78 ± 0.01	10.51 ± 0.23	9.03 ± 0.56	10.48 ± 0.23	6.73 ± 1.40	8.41 ± 0.15	8.72 ± 1.07	10.66 ± 0.09
MH-PcGrad	4.84 ± 0.01	5.55 ± 0.11	5.51 ± 0.03	3.68 ± 0.70	4.67 ± 0.10	10.64 ± 0.01	9.64 ± 0.33	10.55 ± 0.12	9.95 ± 1.93	9.02 ± 0.24	9.90 ± 0.25	10.61 ± 0.03
MH-GradNorm	4.63 ± 0.15	5.37 ± 0.17	5.54 ± 0.20	3.28 ± 0.90	3.81 ± 1.20	10.64 ± 0.01	8.86 ± 1.27	10.26 ± 0.28	7.45 ± 0.75	7.87 ± 1.06	8.16 ± 2.21	10.31 ± 0.22
ParetoFlow	4.23 ± 0.04	5.65 ± 0.11	5.29 ± 0.14	5.00 ± 0.22	4.48 ± 0.11	10.60 ± 0.02	10.13 ± 0.16	10.41 ± 0.09	10.29 ± 0.17	9.65 ± 0.23	9.25 ± 0.43	8.94 ± 0.18
PGD-MOO	4.41 ± 0.08	5.33 ± 0.05	5.54 ± 0.10	5.02 ± 0.03	**4.82 ± 0.01**	10.65 ± 0.01	10.55 ± 0.01	**10.63 ± 0.01**	10.64 ± 0.01	10.06 ± 0.02	**10.14 ± 0.01**	9.70 ± 0.18
SPREAD	**4.89 ± 0.02**	**6.52 ± 0.00**	5.82 ± 0.04	4.90 ± 0.13	4.51 ± 0.04	**11.46 ± 0.13**	**13.27 ± 0.02**	10.23 ± 0.01	**31.19 ± 1.10**	**10.85 ± 0.00**	9.56 ± 0.42	**11.35 ± 0.00**

HV (↑)	$m=2$			$m=3$						$m=4$		$m=6$
Method	RE21	RE22	RE25	RE31	RE32	RE33	RE35	RE36	RE37	RE41	RE42	RE61
Ȳ(best)	4.10	4.78	4.79	10.6	10.56	10.56	10.08	7.61	5.57	18.27	14.52	97.49
MM	4.60 ± 0.00	4.84 ± 0.00	4.63 ± 0.25	10.65 ± 0.00	10.62 ± 0.02	10.62 ± 0.00	10.55 ± 0.01	10.24 ± 0.03	6.73 ± 0.03	20.77 ± 0.08	22.59 ± 0.11	108.96 ± 0.06
MM-COM	4.38 ± 0.09	4.84 ± 0.00	4.83 ± 0.01	10.64 ± 0.01	10.64 ± 0.01	10.61 ± 0.00	10.55 ± 0.02	9.82 ± 0.35	6.35 ± 0.10	20.37 ± 0.06	17.44 ± 0.71	107.99 ± 0.48
MM-IOM	4.58 ± 0.02	4.84 ± 0.00	4.83 ± 0.01	10.65 ± 0.00	10.65 ± 0.00	10.62 ± 0.00	10.57 ± 0.01	**10.29 ± 0.04**	6.71 ± 0.02	20.66 ± 0.05	22.43 ± 0.10	107.71 ± 0.50
MM-ICT	4.60 ± 0.00	4.84 ± 0.00	4.84 ± 0.00	10.65 ± 0.00	10.64 ± 0.00	10.62 ± 0.00	10.50 ± 0.01	10.29 ± 0.03	6.73 ± 0.00	20.58 ± 0.04	22.27 ± 0.15	108.68 ± 0.27
MM-RoMA	4.57 ± 0.00	4.61 ± 0.51	4.83 ± 0.01	10.64 ± 0.01	10.64 ± 0.00	10.58 ± 0.03	10.53 ± 0.03	9.72 ± 0.28	6.67 ± 0.02	20.39 ± 0.09	21.41 ± 0.37	108.47 ± 0.28
MM-TriMentoring	4.60 ± 0.00	4.84 ± 0.00	4.84 ± 0.00	10.65 ± 0.00	10.62 ± 0.01	10.60 ± 0.01	10.59 ± 0.00	9.64 ± 1.42	6.73 ± 0.01	20.68 ± 0.04	21.60 ± 0.19	108.61 ± 0.29
MH	4.60 ± 0.00	4.84 ± 0.00	4.74 ± 0.20	10.65 ± 0.00	10.60 ± 0.05	10.62 ± 0.00	10.49 ± 0.07	10.23 ± 0.03	6.67 ± 0.05	20.62 ± 0.11	22.38 ± 0.35	108.92 ± 0.22
MH-PcGrad	4.59 ± 0.01	4.73 ± 0.36	4.78 ± 0.14	10.64 ± 0.01	10.63 ± 0.01	10.59 ± 0.03	10.51 ± 0.05	10.17 ± 0.08	6.68 ± 0.06	20.66 ± 0.10	22.57 ± 0.26	108.54 ± 0.23
MH-GradNorm	4.28 ± 0.39	4.70 ± 0.44	4.52 ± 0.50	10.60 ± 0.10	10.54 ± 0.15	10.03 ± 1.50	9.76 ± 1.30	9.67 ± 0.43	5.67 ± 1.41	17.06 ± 3.82	18.77 ± 2.99	108.01 ± 1.00
ParetoFlow	4.20 ± 0.17	4.86 ± 0.01	4.84 ± 0.00	10.66 ± 0.12	10.61 ± 0.00	10.75 ± 0.20	**11.12 ± 0.02**	8.42 ± 0.35	6.55 ± 0.59	19.41 ± 0.92	20.35 ± 5.31	107.10 ± 6.96
PGD-MOO	4.46 ± 0.03	4.84 ± 0.00	4.84 ± 0.00	10.65 ± 0.00	10.51 ± 0.04	10.32 ± 0.10	9.37 ± 0.17	6.13 ± 0.12	19.31 ± 0.46	19.01 ± 0.68	105.02 ± 1.14
SPREAD	**4.83 ± 0.00**	**5.18 ± 0.00**	**5.33 ± 0.06**	**11.87 ± 0.00**	**11.27 ± 0.17**	**13.50 ± 0.02**	10.78 ± 0.03	9.55 ± 0.24	**8.37 ± 0.06**	**22.29 ± 0.26**	**23.18 ± 0.96**	**251.81 ± 49.97**

Bayesian

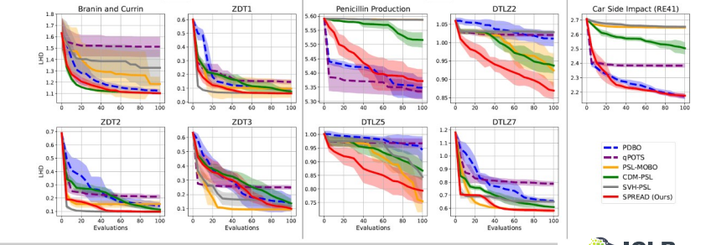

Acknowledgement: This project received funding from the German Federal Ministry of Education and Research (BMBF) through the AI junior research group "Multicriteria Machine Learning".

ICLR

pip install moospread