Breaking Multi-Task Curse: Reward-Weighted Evolution for Black-Box Many-Task Optimization

Yanchi Li, Jiao Liu, Wenyin Gong, Qiong Gu, Yue Zhao, Yew-Soon Ong

ICML 26

MES-RET – Background & Motivation

Black-Box Many-Task Optimization

$$\boldsymbol{x}_i^* = \arg\min f_i(\boldsymbol{x}_i), \ \ \boldsymbol{x}_i \in \mathbb{R}^{D_i},$$

Massive Task Emergence: *A **large number of optimization tasks** appear concurrently*

(e.g., cloud computing scenarios)

Derivative-Free Nature: *First- or second-order **gradient information is inaccessible** for any task*

Applicable Solvers: ***Evolutionary Algorithms**, Bayesian Optimization, etc.*

Our Focus: *Black-box many-task optimization scenarios with **moderate evaluation budget***

MES-RET – Background & Motivation

Evolutionary Multi-Tasking (EMT) *accelerates black-box optimization via* ***Knowledge Transfer*** *among tasks.*
However, it struggles when scaled to ***Many Low-Similarity Tasks.***

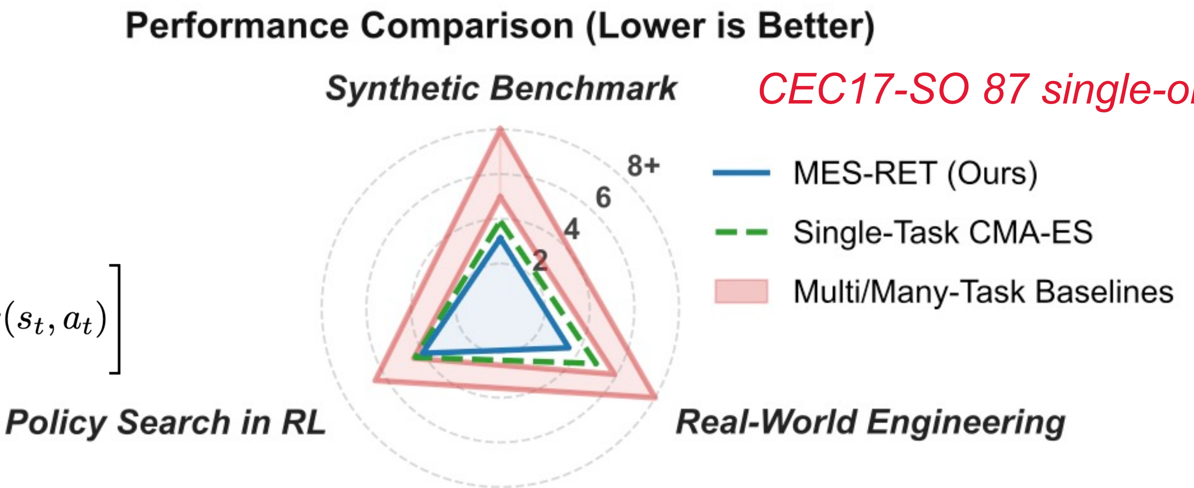

CEC17-SO 87 single-objective tasks in one problem

$$J(\boldsymbol{\theta}) = \mathbb{E}_{\tau\sim\pi_{\boldsymbol{\theta}}}\left[\sum_{t=0}^{T}\gamma^t r(s_t, a_t)\right]$$

Not Traditional High-Similarity
Many-Task Benchmark

Gymnasium 18 tasks in one problem

CEC20-RWCO 42 real-world
constrained tasks in one problem

Empirically, ***Multi-Task Methods Underperformed Single-Task Baseline*** *in such multi-task curse scenarios!*

MES-RET – Background & Motivation

Why EMT Fails at Scale? Multi-Task Curse:

Taylor expansion of expected improvement in CMA-ES $J_k(\boldsymbol{m}_k, \boldsymbol{C}_k) = \mathbb{E}_{\boldsymbol{x}\sim\mathcal{N}}[f_k(\boldsymbol{x})]$

$$\mathcal{I}_{MT} - \mathcal{I}_{ST} \approx \underbrace{\sum_{k=1}^{K} \frac{\partial \mathbb{E}[\Delta_k]}{\partial b_k} \left(b_k - \frac{B}{K} \right)}_{\text{Resource Dilution}} + \underbrace{\sum_{k=1}^{K} \left(\langle \nabla_{\boldsymbol{m}} J_k, \delta \boldsymbol{m}_k \rangle + \mathrm{Tr}(\nabla_{\boldsymbol{C}} J_k \delta \boldsymbol{C}_k) \right)}_{\text{Negative Transfer}},$$

1. Resource Dilution: *Uniform budget allocation forces the first term to be ≤ 0 as tasks scale.*

2.1. Mean Shift: *Conflicting transferred updates oppose the true local gradients.*

2.2. Covariance Distortion: *Maladaptive variance inflates the condition number of the search distribution, driving trace < 0.*

MES-RET – Proposed Method

We Propose ***MES-RET***

Many-task Evolution Strategy *with* ***Reward-weighted Evaluation and Transfer***

For Breaking Multi-Task Curse and for Black-Box Many-Task Optimization

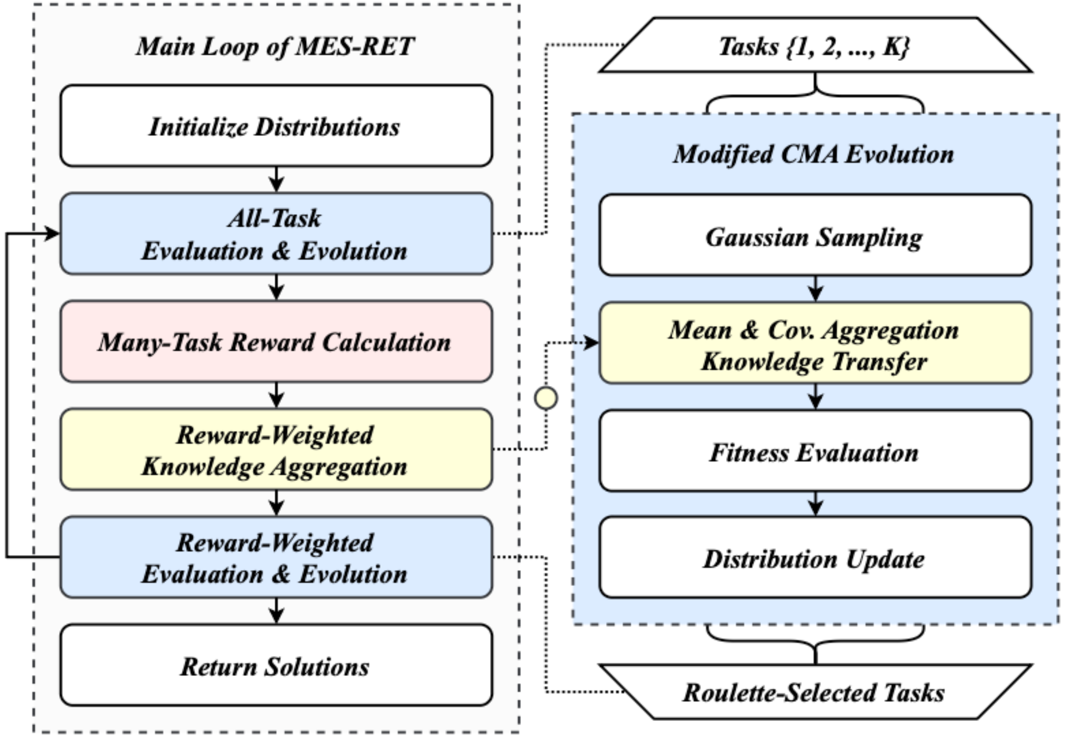

Reward-weighted Evaluation

Reward-weighted Knowledge Transfer

Policy Semantic Alignment

MES-RET – Reward-Weighted Evaluation

Evaluation Budget Allocation

*Replace uniform allocation with **Potential-Driven** scheme:*

Phase 1: All-Task Evolution

Every task gets a baseline budget for state estimation.

Phase 2: Categorical Sampling

High-potential tasks receive evaluations based on reward.

Proposition 3.1 (Effectiveness Condition). *Let r_k be the assigned reward and Δ_k the expected improvement for task k. If the allocation satisfies the positive correlation condition:*

$$K \sum_k r_k \Delta_k \geq (\sum_k r_k)(\sum_k \Delta_k), \tag{4}$$

then the expected total improvement of the reward-weighted strategy exceeds or equals that of the uniform strategy: $\mathbb{E}[\Delta_{total}^{weighted}] \geq \mathbb{E}[\Delta_{total}^{uniform}]$.

Remark: Proposition 3.1 establishes *positive correlation* as the critical prerequisite. This theoretical boundary defines our design objective: the reward signal r_k must serve as a robust proxy for the expected improvement $\mathbb{E}[\Delta_k]$.

Algorithm 1 Reward-Weighted Evaluation

1: **for** $k = 1$ to K **do**
2: Evolve task k: $\mathcal{N}_k \leftarrow$ CMA-Evolution$(\mathcal{N}_k, \lambda)$
3: **end for**
4: $\{r_k\}_{k=1}^K \leftarrow$ Calculate-Reward$(\cdot)$ {See Sec. 3.2}
5: $\{p_k\}_{k=1}^K \leftarrow \{r_k / \sum_{j=1}^K r_j\}_{k=1}^K$
6: **for** 1 to K **do**
7: Select task: $t \leftarrow$ Categorical-Sampling$(\{p_k\}_{k=1}^K)$
8: Evolve task t: $\mathcal{N}_t \leftarrow$ CMA-Evolution$(\mathcal{N}_t, \lambda)$
9: **end for**

Reward 1: Performance

Historical progress reflects future local improvement trends.

$$r_k^{\text{perf}} = \frac{\exp(\bar{\mathcal{I}}_k)}{\sum_j \exp(\bar{\mathcal{I}}_j)}, \quad \mathcal{I}_k = \frac{f_k^{\text{pre}} - f_k^{\text{cur}}}{|f_k^{\text{init}} - f_k^{\text{pre}}| + \epsilon},$$

Reward 2: Diversity

Entropy retention preserves global search potential.

$$r_k^{\text{div}} = \frac{\sigma_k \cdot \text{trace}(\boldsymbol{C}_k)}{n_k},$$

Dynamic Reward Composition

$$r_k = \begin{cases} r_k^{\text{perf}} & \text{if } u \geq \rho, \quad u \sim \mathcal{U}(0,1), \\ r_k^{\text{div}} & \text{otherwise.} \end{cases}$$

MES-RET – Reward-Weighted Knowledge Transfer

Inter-Task Knowledge Transfer

Guiding Search via **Reward-weighted Aggregation**:

Mean Aggregation:

*Pools the directional vectors of high-reward tasks to guide search toward universally **promising regions**.*

Covariance Aggregation:

*Transfers successful **adaptation rhythms** without disrupting a task's specific local correlation structure.*

The Safety Valve - **Weak Guidance Injection**:

Only a marginal fraction of external solutions are injected.

Proposition 3.3 (Mitigation of Negative Transfer). *Let the distribution parameters be updated via the weak guidance injection (Definition 3.2), where a set of external solutions $\mathcal{S}_{trans} = \{\bar{\boldsymbol{x}}_i\}_{i=1}^{2\tau}$ replaces a subset of native candidates. Let $\mathcal{S}_{native}^{\mu}$ denote the set of the top-μ elite solutions derived purely from the native distribution. If the injected solutions induce inferior performance such that $\min_{\bar{\boldsymbol{x}} \in \mathcal{S}_{trans}} \mathcal{L}(\bar{\boldsymbol{x}}) > \max_{\boldsymbol{x} \in \mathcal{S}_{native}^{\mu}} \mathcal{L}(\boldsymbol{x})$, then the distribution parameter update $\Delta\theta$ is explicitly decoupled from $\mathcal{S}_{trans}$. Consequently, the algorithm guarantees the mitigation of negative transfer by assigning zero weight to all solutions in $\mathcal{S}_{trans}$.*

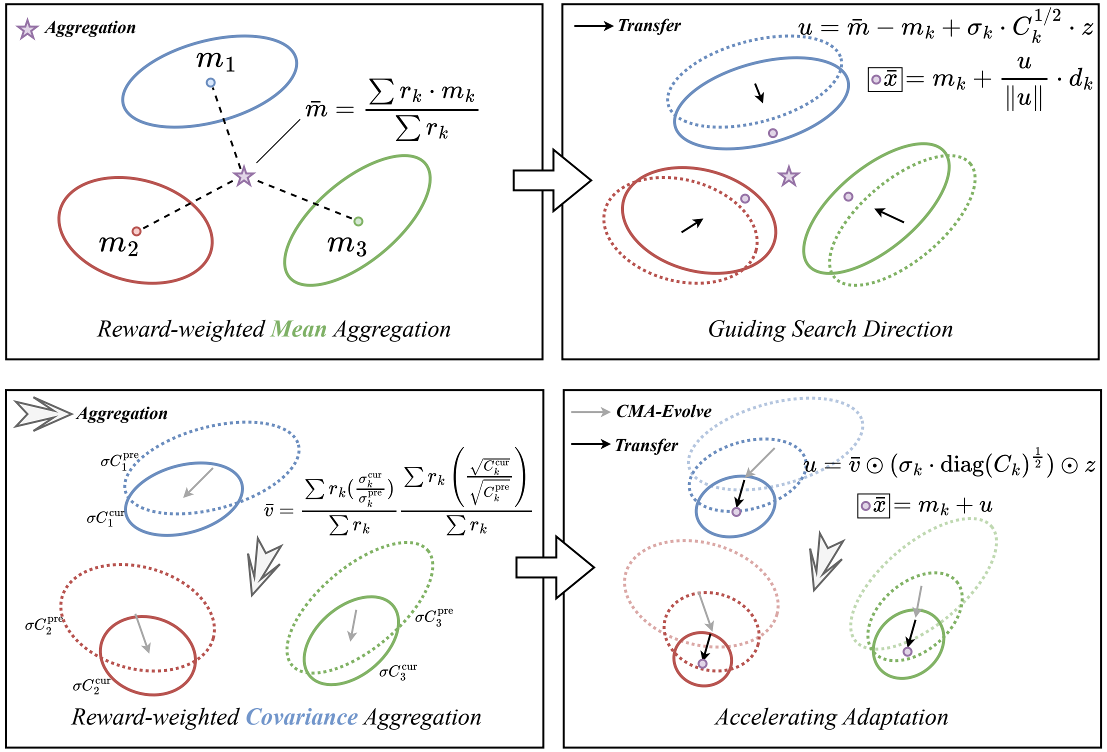

Remark: Proposition 3.3 formalizes a *safety valve* mechanism: external solutions undergo competitive ranking and are automatically discarded if maladaptive. This ensures that the algorithm seamlessly reverts to the standard CMA-ES behavior, effectively neutralizing negative transfer risks.

MES-RET – Policy Semantic Alignment

Policy Parameter Semantic Alignment

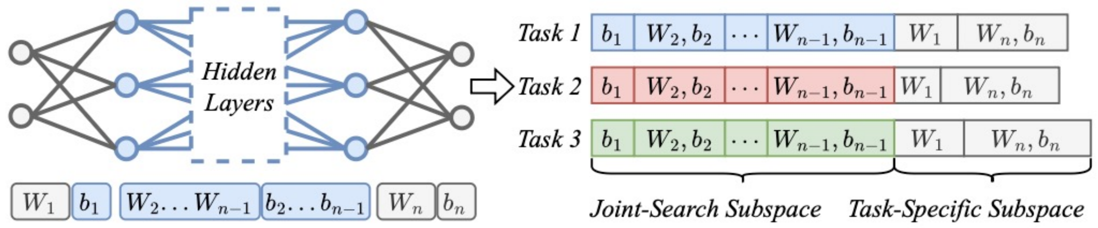

Figure 5. Policy parameter alignment for MLPs. The *joint-search* subspace maintains alignment for inter-task transfer, while the *task-specific* subspace accommodates varying input/output dimensions.

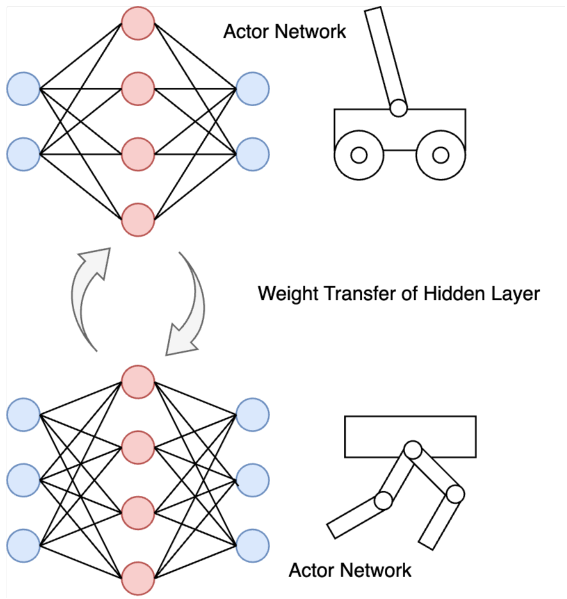

Joint-Search Subspace:

*The shared hidden layers encode **generalized dynamic logic** and transferable motor control primitive.*

Task-Specific Subspace:

*The Input/Output layers that act as **isolated sensor-motor interfaces**, uniquely evolved for each environment's.*

MES-RET – *Experimental Setup*

Not Traditional High-Similarity Many-Task Benchmark

Three Multi-Task Curse Scenarios

***87 Synthetic**, **42 Real-world Engineering**, and **18 Gymnasium RL** tasks.*

CEC17-SO

Task	Description	n
	Unimodal Functions	
T1, T30, T59	Shifted and Rotated Bent Cigar	10, 30, 50
T2, T31, T60	Shifted and Rotated Zakharov	10, 30, 50
	Multimodal Functions	
T3, T32, T61	Shifted and Rotated Rosenbrock	10, 30, 50
T4, T33, T62	Shifted and Rotated Rastrigin	10, 30, 50
T5, T34, T63	Shifted and Rotated Expanded Scaffer F6	10, 30, 50
T6, T35, T64	Shifted and Rotated Lunacek Bi-Rastrigin	10, 30, 50
T7, T36, T65	Shifted and Rotated Non-Continuous Rastrigin	10, 30, 50
T8, T37, T66	Shifted and Rotated Levy	10, 30, 50
T9, T38, T67	Shifted and Rotated Schwefel	10, 30, 50
	Hybrid Functions: $F = f_{1,[1,d_1]} + f_{2,[d_1,d_2]} + \dots$	
T10, T39, T68	Hy1: Zakharov, Rosenbrock, Rastrigin	10, 30, 50
T11, T40, T69	Hy2: Elliptic, Schwefel, Bent Cigar	10, 30, 50
T12, T41, T70	Hy3: Bent Cigar, Rosenbrock, Lunacek Bi-Rastrigin	10, 30, 50
T13, T42, T71	Hy4: Elliptic, Ackley, Schaffer F6, Rastrigin	10, 30, 50
T14, T43, T72	Hy5: Bent Cigar, HGBat, Rastrigin, Rosenbrock	10, 30, 50
T15, T44, T73	Hy6: Schaffer F6, HGBat, Rosenbrock, Schwefel	10, 30, 50
T16, T45, T74	Hy7: Katsuura, Ackley, Griewank, Schwefel, Rastrigin	10, 30, 50
T17, T46, T75	Hy8: Elliptic, Ackley, Rastrigin, HGBat, Discus	10, 30, 50
T18, T47, T76	Hy9: Bent Cigar, Rastrigin, Griewank, Weierstrass, Schaffer F6	10, 30, 50
T19, T48, T77	Hy10: Happycat, Katsuura, Ackley, Rastrigin, Schwefel, Schaffer F6	10, 30, 50
	Composition Functions: $F = w_1 f_1 + w_2 f_2 + \dots$	
T20, T49, T78	Co1: Rosenbrock, Elliptic, Rastrigin	10, 30, 50
T21, T50, T79	Co2: Rastrigin, Griewank, Schwefel	10, 30, 50
T22, T51, T80	Co3: Rosenbrock, Ackley, Schwefel, Rastrigin	10, 30, 50
T23, T52, T81	Co4: Ackley, Elliptic, Griewank, Rastrigin	10, 30, 50
T24, T53, T82	Co5: Rastrigin, Happycat, Ackley, Discus, Rosenbrock	10, 30, 50
T25, T54, T83	Co6: Schaffer F6, Schwefel, Griewank, Rosenbrock, Rastrigin	10, 30, 50
T26, T55, T84	Co7: HGBat, Rastrigin, Schwefel, Bent Cigar, Elliptic, Schaffer F6	10, 30, 50
T27, T56, T85	Co8: Ackley, Griewank, Discus, Rosenbrock, HappyCat, Schaffer F6	10, 30, 50
T28, T57, T86	Co9: Hy5, Hy6, Hy7	10, 30, 50
T29, T58, T87	Co10: Hy5, Hy8, Hy9	10, 30, 50

CEC20-RWCO

Task	Description	n	g	h
	Industrial chemical processes			
T1	Heat exchanger network design (case 1)	9	0	8
T2	Heat exchanger network design (case 2)	11	0	9
T3	Optimal operation of Alkylation unit	7	14	0
T4	Reactor network design	6	1	4
T5	Haverly's pooling problem	9	2	4
	Process synthesis and design problems			
T6	Process synthesis problem	2	2	0
T7	Process synthesis and design problem	3	1	1
T8	Process flow sheeting problem	3	3	0
T9	Two-reactor problem	7	4	4
T10	Process synthesis problem	7	9	0
T11	Process design problem	5	3	0
T12	Multi-product batch plant	10	10	0
	Mechanical engineering problem			
T13	Weight minimization of a speed reducer	7	11	0
T14	Optimal design of industrial refrigeration system	14	15	0
T15	Tension/compression spring design (case 1)	3	3	0
T16	Pressure vessel design	4	4	0
T17	Welded beam design	4	5	0
T18	Three-bar truss design problem	2	3	0
T19	Multiple disk clutch brake design problem	5	7	0
T20	Planetary gear train design optimization problem	9	10	1
T21	Step-cone pulley problem	5	8	3
T22	Robot gripper problem	7	7	0
T23	Hydro-static thrust bearing design problem	4	7	0
T24	Four-stage gearbox problem	22	86	0
T25	10-bar truss design	10	3	0
T26	Rolling element bearing	10	9	0
T27	Gas transmission compressor design	4	1	0
T28	Tension/compression spring design (case 2)	3	8	0
T29	Gear train design Problem	4	1	1
T30	Himmelblau's function	5	6	0
T31	Topology optimization	30	30	0
	Power Electronic Problems			
T32	Wind farm layout problem	30	91	0
T33	Synchronous optimal pulse-width modulation (SOPM) for 3-level invereters	25	24	1
T34	SOPM for 5-level inverters	25	24	1
T35	SOPM for 7-level inverters	25	24	1
T36	SOPM for 9-level inverters	30	29	1
T37	SOPM for 11-level inverters	30	29	1
T38	SOPM for 13-level inverters	30	29	1
	Livestock Feed Ration Optimization			
T39	Beef Cattle (case 1)	59	14	1
T40	Beef Cattle (case 2)	59	14	1
T41	Beef Cattle (case 3)	59	14	1
T42	Beef Cattle (case 4)	59	14	1

Gymnasium policy search

2 hidden-layer MLP, each with 16 neurons

Task	Description	Observation		Action	
	Classical Control				
T1	MountainCarContinuous-v0	2	Con.	1	Con.
T2	MountainCar-v0	2	Con.	3	Dis.
T3	Pendulum-v1	3	Con.	1	Con.
T4	CartPole-v1	4	Con.	2	Dis.
T5	Acrobot-v1	6	Con.	3	Dis.
	Box2D				
T6	LunarLander-v3	8	Con.	4	Dis.
T7	BipedalWalker-v3	24	Con.	4	Con.
	MuJoCo				
T8	InvertedPendulum-v5	4	Con.	1	Con.
T9	InvertedDoublePendulum-v5	9	Con.	1	Con.
T10	Reacher-v5	10	Con.	2	Con.
T11	Pusher-v5	23	Con.	7	Con.
T12	HalfCheetah-v5	17	Con.	6	Con.
T13	Hopper-v5	11	Con.	3	Con.
T14	Walker2d-v5	17	Con.	6	Con.
T15	Swimmer-v5	8	Con.	2	Con.
T16	Ant-v5	105	Con.	8	Con.
T17	Humanoid-v5	348	Con.	17	Con.
T18	HumanoidStandup-v5	378	Con.	17	Con.

MES-RET – ***Experimental Setup***

Comparison Baselines

Evolution Strategies, ***Population-based EAs***, *and* ***Policy-gradient Methods***

Algorithm	Description	Task	SO	CO	PS	Type
CMA-ES (Hansen & Ostermeier, 2001)	Covariance Matrix Adaptation Evolution Strategy	Single	✓	✓	✓	Evolution Strategy
OpenAI-ES (Salimans et al., 2017)	OpenAI Evolution Strategy with Stochastic Gradient Descent	Single			✓	
SBCMAES (Liaw & Ting, 2019)	Many-Task CMA-ES based on Symbiosis in Biocoenosis	Many	✓	✓	✓	
MTES-KG (Li et al., 2024e)	Multi-Task CMA-ES with Knowledge-Guided External Sampling	Multi	✓	✓	✓	
TNG-SNES (Li et al., 2024c)	Many-Task Natural ES with Task-Averaged Natural Gradient Transfer	Many	✓	✓	✓	
L-SHADE (Tanabe & Fukunaga, 2014)	Success History based Adaptive DE with Linear Population Reduction	Single	✓			Differential Evolution
CCEF-ECHT (Li et al., 2024d)	Competitive and Cooperative DE for Constrained Optimization	Single		✓		
CEDA-MP (Zhang et al., 2024)	Constrained Multi-Task DE via Co-Evolution and Domain Adaptation	Multi		✓		
AT-MFEA (Xue et al., 2022)	Affine Transformation Enhanced Heterogeneous Multi-Task GA	Multi	✓	✓		Genetic Algorithm
EMaTO-MKT (Liang et al., 2022)	Many-Task GA based on Multi-Source Knowledge Transfer	Many	✓			
PPO (Schulman et al., 2017)	Proximal Policy Optimization with Clipped Surrogate Objective	Single			✓	Policy Gradient
A2C (Mnih et al., 2016)	Advantage Actor-Critic with Generalized Advantage Estimation	Single			✓	

MES-RET – Experimental Results

Comparison & Ablation Results

*MES-RET achieves the **lowest overall Friedman Rankings (3.13, 3.54, 4.03)** and effectively outperforms single-task baselines where other multi-task methods fail.*

Synthetic Optimization (87 tasks)						Real-World Application (42 tasks)					
Algorithm	+	−	=	#B	Rank	Algorithm	+	−	=	#B	Rank
MES-RET	↓	↑	/	**25**	**3.13**	**MES-RET**	↓	↑	/	**18**	**3.54**
w/o-RE	1	17	69	19	3.82	w/o-RE	4	11	27	15	3.65
w/o-RT	1	7	79	23	3.28	w/o-RT	5	10	27	17	3.79
CMA-ES	6	21	60	20	3.90	CMA-ES	7	19	16	15	4.95
SBCMAES	1	83	3	0	9.68	SBCMAES	1	35	6	2	8.90
MTES-KG	4	43	40	13	4.97	MTES-KG	2	31	9	3	6.75
TNG-SNES	7	63	17	7	6.70	TNG-SNES	0	30	12	5	7.05
L-SHADE	20	57	10	19	5.33	CCEF-ECHT	9	17	16	17	4.23
AT-MFEA	8	72	7	4	7.02	AT-MFEA	5	29	8	4	6.24
EMaTO-MKT	5	73	9	4	7.17	CEDA-MP	4	28	10	4	5.90

Gymnasium 18 Tasks

Algorithm	+	−	=	#B	Rank
sep-MES-RET	↓	↑	/	**6**	**4.03**
w/o-RE	1	3	14	3	5.42
w/o-RT	3	6	9	4	5.08
sep-CMA-ES	1	5	12	3	4.39
sep-SBCMAES	2	7	9	3	6.47
sep-MTES-KG	0	7	11	2	6.31
TNG-SNES	1	5	12	4	4.47
OpenAI-ES	0	4	14	3	4.67
PPO	3	10	5	**6**	5.75
A2C	0	14	4	1	8.42

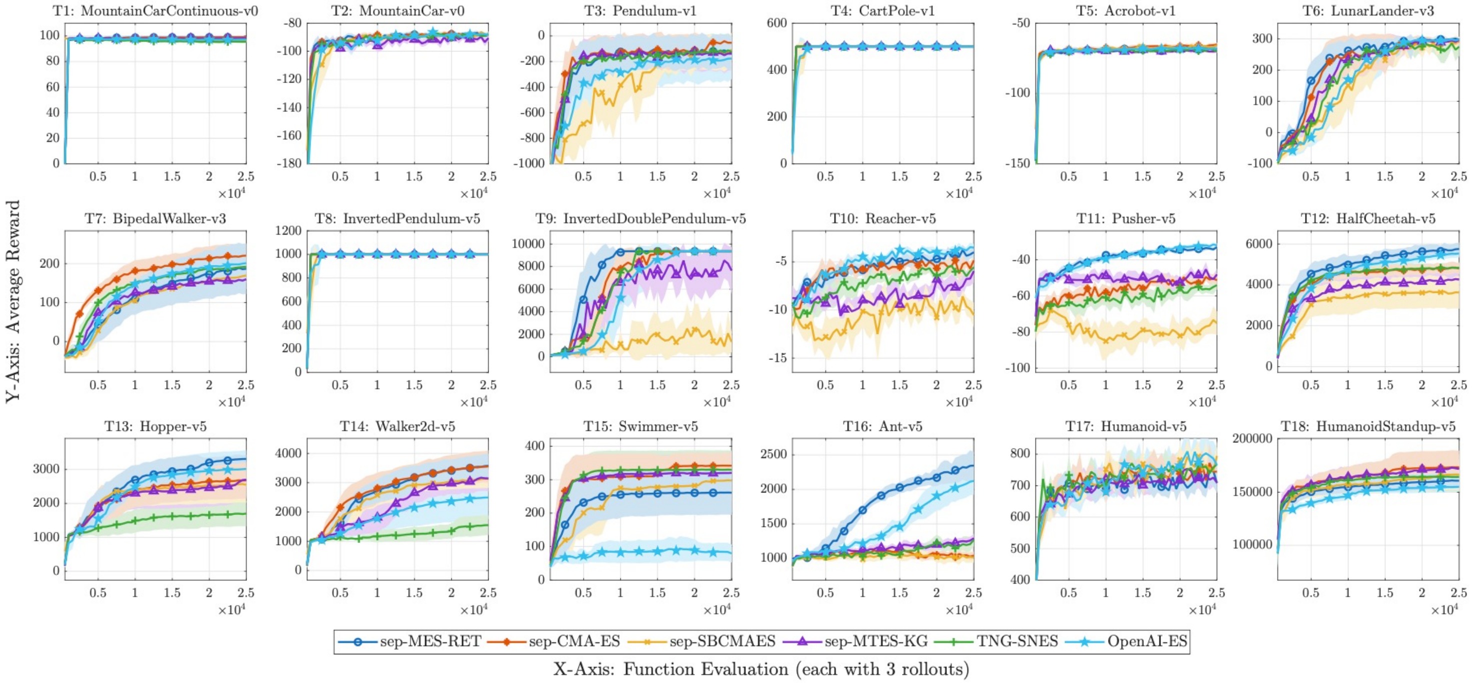

MES-RET – Experimental Results

Behavioral Skill Transfer Insights:

Parameter Correlation:

high cross-task correlation within the joint-search subspace.

Emergent Consistency:

Morphologically distinct agents learned highly consistent forward-propulsion mechanisms.

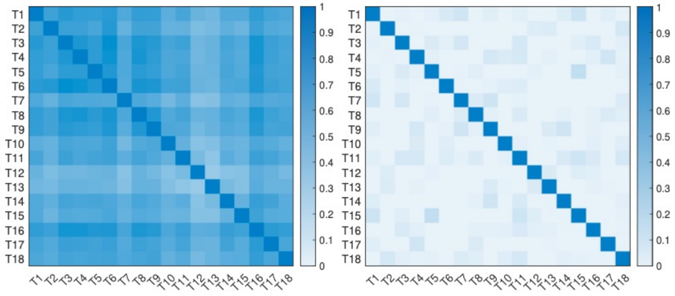

(a) sep-MES-RET *(b)* w/o-RT

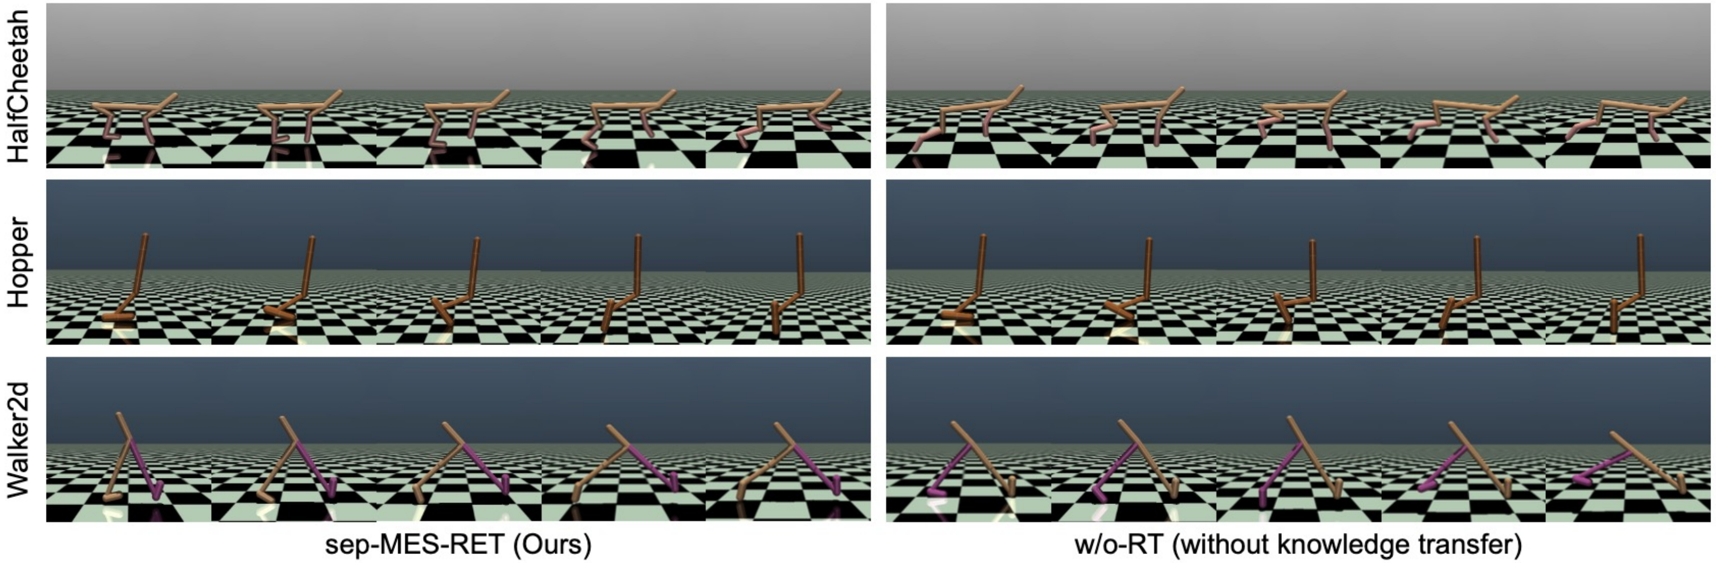

Figure 7. Visualization of the learned locomotion policies by proposed sep-MES-RET (left) and w/o-RT (right) across HalfCheetah, Hopper, and Walker2d environments. sep-MES-RET discovered transferable movement patterns: all three agents learned similar forward propulsion mechanisms using synchronized leg extension and kicking motions. In contrast, w/o-RT, without joint-search subspace based knowledge transfer, developed inconsistent and distinct movement strategies for each environment.

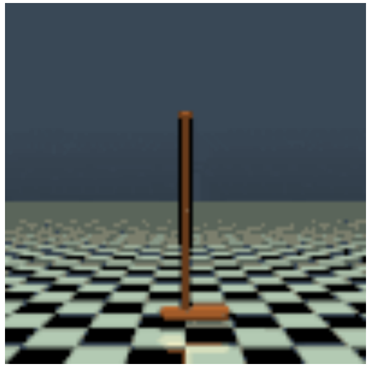

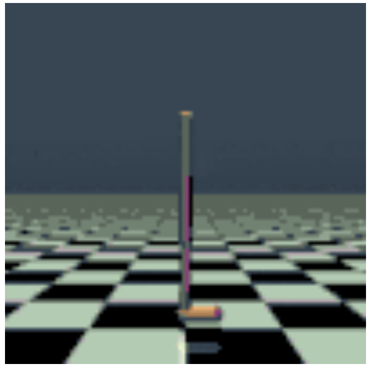

MToP – Benchmarking Platform for Evolutionary Multitasking

- **100+** multi/single-task evolutionary algorithms
- **150+** benchmark problems with applications
- **20+** performance metrics for evaluation
- **3** non-parametric statistical analysis method
- **M**ultiple optimization types with combinations
- **V**arious visualization functions and tools
- **G**lobal random seed control

https://github.com/intLyc/MTO-Platform

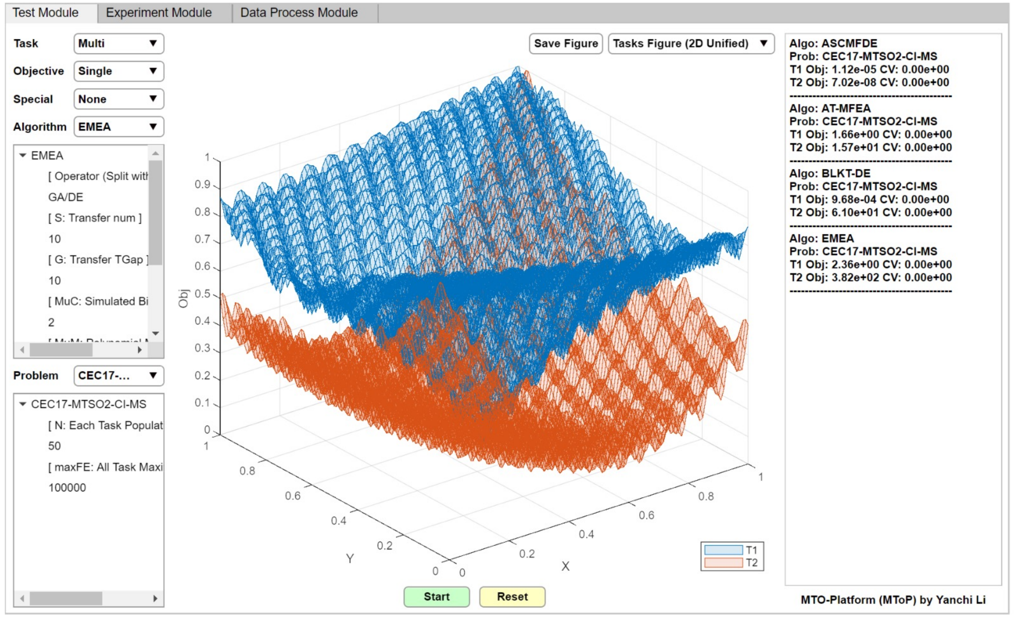

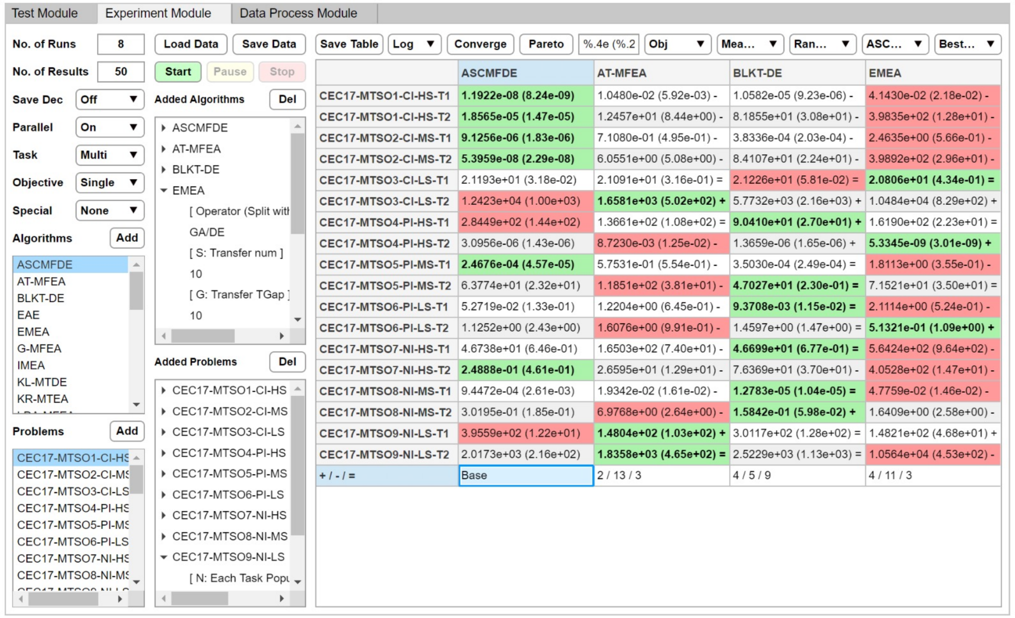

	ASCMFDE	AT-MFEA	BLKT-DE	EMEA
CEC17-MTSO1-CI-HS-T1	**1.1922e-08 (8.24e-09)**	1.0480e-02 (5.92e-03) -	1.0582e-05 (9.23e-06) -	4.1430e-02 (2.18e-02) -
CEC17-MTSO1-CI-HS-T2	**1.8565e-05 (1.47e-05)**	1.2457e+01 (8.44e+00) -	8.1855e+01 (3.08e+01) -	3.9835e+02 (1.28e+01) -
CEC17-MTSO2-CI-MS-T1	**9.1256e-06 (1.83e-06)**	7.1080e-01 (4.95e-01) -	3.8336e-04 (2.03e-04) -	2.4635e+00 (5.66e-01) -
CEC17-MTSO2-CI-MS-T2	**5.3959e-08 (2.29e-08)**	6.0551e+00 (5.08e+00) -	8.4107e+01 (2.24e+01) -	3.9892e+02 (2.96e+01) -
CEC17-MTSO3-CI-LS-T1	2.1193e+01 (3.18e-02)	2.1091e+01 (3.16e-01) =	2.1226e+01 (5.81e-02) =	**2.0806e+01 (4.34e-01) =**
CEC17-MTSO3-CI-LS-T2	1.2423e+04 (1.00e+03)	**1.6581e+03 (5.02e+02) +**	5.7732e+03 (2.16e+03) +	1.0484e+04 (8.29e+02) +
CEC17-MTSO4-PI-HS-T1	2.8449e+02 (1.44e+02)	1.3661e+02 (1.08e+02) =	**9.0410e+01 (2.70e+01) +**	1.6190e+02 (2.23e+01) =
CEC17-MTSO4-PI-HS-T2	3.0956e-06 (1.43e-06)	8.7230e-03 (1.25e-02) -	1.3659e-06 (1.65e-06) +	**5.3345e-09 (3.01e-09) +**
CEC17-MTSO5-PI-MS-T1	**2.4676e-04 (4.57e-05)**	5.7531e-01 (5.54e-01) -	3.5030e-04 (2.49e-04) =	1.8113e+00 (3.55e-01) -
CEC17-MTSO5-PI-MS-T2	6.3774e+01 (2.32e+01)	1.1851e+02 (3.81e+01) -	**4.7027e+01 (2.30e-01) =**	7.1521e+01 (3.50e+01) =
CEC17-MTSO6-PI-LS-T1	5.2719e-02 (1.33e-01)	1.2204e+00 (6.45e-01) -	**9.3708e-03 (1.15e-02) =**	2.1114e+00 (5.24e-01) -
CEC17-MTSO6-PI-LS-T2	1.1252e+00 (2.43e+00)	1.6076e+00 (9.91e-01) -	1.4597e+00 (1.47e+00) =	**5.1321e-01 (1.09e+00) +**
CEC17-MTSO7-NI-HS-T1	4.6738e+01 (6.46e-01)	1.6503e+02 (7.40e+01) -	**4.6699e+01 (6.77e-01) =**	5.6424e+02 (9.64e+02) -
CEC17-MTSO7-NI-HS-T2	**2.4888e-01 (4.61e-01)**	2.6595e+01 (1.29e+01) -	7.6369e+01 (3.70e+01) -	4.0528e+02 (1.47e+01) -
CEC17-MTSO8-NI-MS-T1	9.4472e-04 (2.61e-03)	1.9342e-02 (1.61e-02) -	**1.2783e-05 (1.04e-05) =**	4.7759e-02 (1.46e-02) -
CEC17-MTSO8-NI-MS-T2	3.0195e-01 (1.85e-01)	6.9768e+00 (2.64e+00) -	**1.5842e-01 (5.98e-02) +**	1.6409e+00 (2.58e+00) -
CEC17-MTSO9-NI-LS-T1	3.9559e+02 (1.22e+01)	**1.4804e+02 (1.03e+02) +**	3.0117e+02 (1.28e+02) =	1.4821e+02 (4.68e+01) +
CEC17-MTSO9-NI-LS-T2	2.0173e+03 (2.16e+02)	**1.8358e+03 (4.65e+02) =**	2.5229e+03 (1.13e+03) =	1.0564e+04 (4.53e+02) -
+ / - / =	Base	2 / 13 / 3	4 / 5 / 9	4 / 11 / 3

Breaking Multi-Task Curse:
Reward-Weighted Evolution for Black-Box Many-Task Optimization

Thanks!

Yanchi Li, *Jiao Liu, Wenyin Gong, Qiong Gu, Yue Zhao, Yew-Soon Ong*

ICML 26