



HAL
open science

Meta-learning with similarity-weighted transfer for manufacturing optimization

Madani Bezoui, Hajar Nouinou, Ahcène Bounceur

► **To cite this version:**

Madani Bezoui, Hajar Nouinou, Ahcène Bounceur. Meta-learning with similarity-weighted transfer for manufacturing optimization. Société Française de Recherche Opérationnelle et d'Aide à la Décision, Feb 2026, Tours, France. <hal-05548907>

HAL Id: hal-05548907

<https://hal.science/hal-05548907v1>

Submitted on 17 Mar 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Copyright - All rights reserved

Meta-learning with similarity-weighted transfer for manufacturing optimization

Madani Bezoui¹, Hajar Nouinou¹, Ahcene Bounceur²

¹ CESI LINEACT, Nancy, 54000, France
{mbezoui, hnouinou}@cesi.fr

² College of Computing and Informatics, University of Sharjah, United Arab Emirates

Keywords : *Bayesian optimization, Meta-learning, Transfer learning, Manufacturing optimization, Sample efficiency, Industry 4.0*

1 Introduction

In Industry 4.0 and 5.0, manufacturing reconfiguration optimizes production parameters (speed, temperature, pressure, tooling sequences) when adapting lines to new products. These problems exhibit: (i) high-dimensional search spaces (5 to 15 parameters), (ii) expensive black-box evaluations via physical experiments (€50 to 500 per trial), (iii) multiobjective trade-offs (throughput, quality, energy), and (iv) strict budgets of 20 to 50 evaluations due to downtime and material costs. Classical Design of Experiments suffers combinatorial explosion, requiring hundreds of evaluations. Evolutionary algorithms, while balancing exploration-exploitation through mutation and crossover, require 500 to 5,000 evaluations—exceeding manufacturing budgets by one to two orders of magnitude—and lack uncertainty quantification for risk assessment.

Bayesian Optimization addresses these limitations through probabilistic surrogates and principled exploration-exploitation. However, standard BO faces: the *cold-start problem* wasting initial evaluations despite historical data, uniform search treatment ignoring operational preferences, and independent task optimization discarding shared knowledge. MBAE (Meta-Bayesian Active Explorer) addresses these via hierarchical Bayesian meta-learning with similarity-weighted transfer designed for manufacturing constraints and reconfigurable systems [1].

2 Framework architecture

MBAE integrates three innovations building on meta-learning and transfer learning [3]:

Hierarchical Bayesian meta-learning: Standard BO starts with uninformative priors, wasting early evaluations. MBAE learns hyperparameter priors $\theta_i \sim \mathcal{N}(\mu_0, \Sigma_0)$ from historical tasks, where θ_i includes length-scale and noise parameters of the Gaussian Process (GP) surrogate model. The hierarchical structure models μ_0 and Σ_0 as global meta-parameters inferred from past optimization campaigns, enabling informed initialization while preserving uncertainty for risk-aware manufacturing decisions.

Similarity-weighted transfer: To prevent negative transfer from unrelated tasks, MBAE computes task similarity using a learned feature mapping $\phi(\cdot)$ that embeds tasks into a metric space based on problem characteristics (dimensionality, smoothness, constraint structure). The similarity weight $z(T_i, T_{\text{new}}) = \exp\left(-\frac{\|\phi(T_i) - \phi(T_{\text{new}})\|^2}{2\ell_\phi^2}\right)$, down-weights contributions from dissimilar tasks, where ℓ_ϕ is a bandwidth parameter controlling sensitivity, and $\phi(T_i)$ represents the feature vector of task T_i . Entropy-based scaling further attenuates unreliable transfers.

Manufacturing-aware acquisition: Following GP-UCB principles [2], MBAE uses a composite acquisition function $\alpha_{\text{MBAE}}(x) = \lambda_1\sigma_t(x) + \lambda_2\text{EI}(x) + \lambda_3w(x)$, where $\sigma_t(x)$ is the GP predictive standard deviation (exploration term), $\text{EI}(x)$ is the Expected Improvement over the current best (exploitation term), and $w(x)$ encodes operational preferences (e.g., safety margins, preferred operating ranges). The adaptive weights $\lambda_i(t)$ shift from exploration to exploitation over time, addressing manufacturing constraints that evolutionary algorithms do not explicitly model through probabilistic uncertainty.

3 Theoretical guarantees

Building on the regret analysis framework for Gaussian process optimization [2], we establish conditional convergence guarantees under bounded prior mismatch. We demonstrate regret bounds $R_T \leq C_1\sqrt{T\gamma_T} + C_2T^{1/2} + C_3 \cdot \Delta\mu \cdot T$ where γ_T represents the maximum information gain and the linear term $C_3\Delta\mu T$ captures cumulative prior mismatch effects from meta-learned initialization. For well-matched tasks with $\Delta\mu \rightarrow 0$, MBAE achieves sublinear regret with multiplicative acceleration factor $\rho(s_{\min}, \Delta\mu)$ over random initialization, providing formal guarantees that translate theoretical advantages into practical performance gains for reconfigurable manufacturing systems where rapid adaptation is critical.

4 Experimental validation and conclusion

Experiments on 9D manufacturing problems (100 runs, 14 algorithms) validate theoretical results with overwhelming significance. MBAE achieves $4.3\times$ to $4.6\times$ convergence speedup, requiring only 6.16 iterations versus 26.6 to 28.3 for baselines to reach 95% optimality (Mann-Whitney U $p < 10^{-4}$, Cohen’s $d = -5.7$ to -6.7). Performance gains include 2.9 to 10.3% over classical BO methods [2], 56.2 to 60.9% over meta-learning baselines [3], and $2.05\times$ computational speedup despite meta-learning overhead.

TAB. 1: Performance comparison

| Method | Iter. | Gain |
|-------------------|-------|--------|
| MBAE | 6.16 | to |
| GP-UCB | 28.3 | +10.3% |
| EI | 26.6 | +2.9% |
| MAML | N/A | +60.9% |
| All $p < 10^{-4}$ | | |

MBAE addresses sample efficiency in expensive manufacturing optimization by intelligently leveraging historical data through principled meta-learning. The framework provides the first integration of hierarchical Bayesian meta-learning with similarity-weighted transfer for Industry 4.0/5.0 reconfigurable systems [1], achieving measurable acceleration with robust negative transfer prevention and theoretical guarantees. Validation across 42,000 evaluations confirms that MBAE translates theoretical advantages into measurable cost savings and reduced timelines for industrial deployment where sample efficiency is paramount.

References

[1] Bezoui, et al. (2023, November). Deep Reinforcement Learning for Multiobjective Scheduling in Industry 5.0 Reconfigurable Manufacturing Systems. In *International Conference on Machine Learning for Networking* (pp. 90-107). Cham: Springer Nature Switzerland.

[2] Srinivas, N., Krause, A., Kakade, S., Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. *ICML*.

[3] Finn, C., Abbeel, P., Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *ICML*.