Table I: System Configuration Parameters

Parameter	Value	Description
Knowledge Graph		
n_topics_initial	15	Initial number of topic clusters
min_cluster_size	10	Minimum chunks per cluster
similarity_threshold	0.7	Threshold for horizontal edges
max_depth	3	Maximum hierarchy depth
embedding_model	all-MiniLM-L6-v2	Sentence transformer model
Graph Statistics		
Total nodes	48	Level 1: 8, Level 2: 23, Level 3: 17
Total edges	156	Hierarchical + similarity connections
Total text chunks	1,247	From 156-page PDF
Differentiable Graph		
embedding_dim	384	Node embedding dimension
Initial experts	5	Selected from largest Level 1 nodes
max_experts	15	Maximum expert capacity
PPO Hyperparameters		
gamma	0.99	Discount factor

Parameter	Value	Description
gae_lambda	0.95	GAE smoothing parameter
clip_epsilon	0.2	PPO clipping threshold
value_coef	0.5	Value loss coefficient
entropy_coef	0.01	Entropy bonus coefficient
ppo_epochs	4	PPO update epochs
batch_size	64	Training batch size
learning_rate	3×10^{-4}	Policy network learning rate
TinyBERT		
lm_learning_rate	2×10^{-5}	Language model learning rate
weight_decay	0.01	L2 regularization
T_0	100	First restart for cosine annealing
T_mult	2	Period multiplier
eta_min	1×10^{-6}	Minimum learning rate
Reward Weights		
perplexity_weight (α)	0.5	Quality component weight
diversity_weight (β)	0.3	Diversity component weight
efficiency_weight (γ)	0.2	Efficiency component weight

Table II: Training Progress at Key Steps

Step	Loss	Reward	Experts	Creations	Deletions	Meta Action
0	10.79	0.1527	1	0	0	-
10	9.82	0.1532	5	4	0	REALLOCATE
20	9.04	0.1541	6	5	0	REALLOCATE
30	8.77	0.1553	8	7	0	REALLOCATE
40	8.43	0.1568	10	9	0	CREATE
50	7.95	0.1584	13	12	0	DELETE
60	7.53	0.1601	15	14	0	REALLOCATE
70	7.26	0.1617	15	14	0	NOOP
80	7.20	0.1632	15	14	0	CREATE
90	6.75	0.1648	15	14	0	DELETE
100	6.26	0.1663	15	14	0	REALLOCATE
110	6.14	0.1678	15	14	0	DELETE
120	6.25	0.1691	15	14	0	NOOP
130	5.95	0.1704	14	15	2	REALLOCATE
140	6.07	0.1718	15	16	2	NOOP
150	5.78	0.1731	15	16	2	NOOP

Step	Loss	Reward	Experts	Creations	Deletions	Meta Action
160	5.73	0.1745	15	16	2	CREATE
170	5.48	0.1758	15	16	2	REALLOCATE
180	5.84	0.1772	15	16	2	NOOP
190	5.76	0.1786	15	16	2	DELETE
200	5.69	0.1801	15	16	2	REALLOCATE
210	5.61	0.1823	15	16	2	REALLOCATE
220	5.84	0.1843	15	16	2	REALLOCATE

Improvement | **-45.8%** | **+20.7%** | **+14** | **+16** | **+2** | - |

Table III: Expert Specialization Metrics

Metric	Value	Interpretation
Global Metrics		
Silhouette Score	0.432	Well-separated clusters (>0.4 = good)
Within-Cluster Distance (mean)	0.34	Low intra-cluster variance
Between-Cluster Distance (mean)	0.61	High inter-cluster separation
t-test p-value	3.2×10^{-8}	Statistically significant separation
Expert Purity	0.78	78% texts correctly assigned
Cluster Composition		
Machining Cluster	5 experts	Turning, milling, grinding specialists
Additive Cluster	5 experts	FDM/SLA, metal AM, DfAM specialists
Materials Cluster	5 experts	Testing, microstructure, thermal specialists
Top Performing Experts		
L1_002_child_4 (Turning)	Loss: 4.92	Best quality (turning operations)
L1_003_child_2 (Additive)	Loss: 5.13	Second best (additive manufacturing)
L1_004_child_1 (Materials)	Loss: 5.28	Third best (materials science)
Worst Performing Experts		
L1_002_child_0 (General)	Loss: 6.87	Candidate for deletion

Metric	Value	Interpretation
L1_003_child_4 (Generic)	Loss: 6.54	Low specialization

Table IV: Knowledge Transfer Analysis

Metric	Value	Description
Transfer Matrix Statistics		
Mean Transfer Strength	0.124	Average off-diagonal transfer
Median Transfer Strength	0.098	50th percentile
Standard Deviation	0.087	Variability in transfer
Maximum Transfer	0.432	L1_002_child_4 → L1_002_child_5
Minimum Transfer	0.012	Between distant clusters
Mean Asymmetry (|T_ij - T_ji|)	0.067	Moderate directional bias
Top 5 Knowledge Transfers		
1. L1_002_child_4 → L1_002_child_5	0.432	Turning → Milling (within cluster)
2. L1_002 → L1_002_child_2	0.398	Parent → Additive child
3. L1_002_child_1 → L1_002_child_3	0.376	Grinding → Surface finishing
4. L1_003 → L1_003_child_0	0.354	Materials → Alloy specialist
5. L1_002_child_5 → L1_002_child_1	0.321	Milling → Grinding
Donor Experts (Net Out > In)		
L1_002 (Parent)	+0.87	Knowledge source for children
L1_002_child_4 (Turning)	+0.63	Specialized donor

Metric	Value	Description
L1_003 (Parent)	+0.51	Materials knowledge source
Recipient Experts (Net In > Out)		
L1_002_child_2 (Additive)	-0.42	Knowledge sink
L1_004_child_0 (CNC)	-0.38	Learning from multiple sources
L1_002_child_5 (Milling)	-0.31	Cross-specialization learner
Transfer Efficiency		
Within-Cluster Transfer	0.287	2.3× stronger than between
Between-Cluster Transfer	0.125	Weaker cross-cluster transfer

Table V: Ablation Study Results

Configuration	Final Loss	Final Reward	Specialization (Silhouette)	Experts	Notes
Full System	**5.84**	**0.1843**	**0.432**	**15**	**Best performance**
No Meta-Controller	7.32	0.167	0.312	5	Population stuck at initial
No Knowledge Transfer	6.91	0.172	0.365	12	Limited specialization
No Diversity Reward	8.45	0.151	0.089	15	Expert collapse (all similar)
Fixed Graph	6.23	0.178	0.398	15	No graph adaptation
Random Actions	9.12	0.149	0.211	8	Uncoordinated behavior
Single Expert	8.76	0.158	N/A	1	No specialization possible

Table VI: Computational Performance

Component	Time per Step	Memory Usage	Notes
Forward Pass			
Knowledge Graph Construction	127 sec (total)	2.1 GB	One-time initialization
Differentiable Graph Forward	0.023 sec	384 MB	Per query
Training (per step)			
Expert Training (15 experts)	4.2 sec	4.8 GB	Parallelizable
Meta-Controller Update	0.34 sec	512 MB	Every 32 steps
Local Policy Updates	1.8 sec	1.2 GB	Every 32 steps
Reward Computation	0.12 sec	256 MB	Per step
Total Training			
220 steps	4.2 hours	5.6 GB peak	NVIDIA T4 GPU
Inference			
Query Routing	0.087 sec	3.2 GB	2-expert response
Response Generation	0.156 sec	3.2 GB	Including TinyBERT