Physics-Informed Few-Shot Learning for Rapid and Data-Efficient Prediction of Manufacturing Quality

A training-free few-shot framework for cross-material relative-density screening in laser powder bed fusion

Prashant Dhakal, Shiren Wang

Department of Industrial & Systems Engineering
Texas A&M University

IISE Annual Conference & Expo 2026

Why LPBF Qualification Hits a Wall

*Metal AM markets are growing fast — but every **new alloy** and every **parameter change** restarts the same multi-year qualification cycle.*

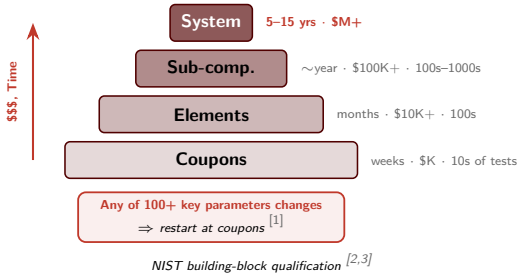

NIST building-block qualification [2,3]

Why LPBF Qualification Hits a Wall

*Metal AM markets are growing fast — but every **new alloy** and every **parameter change** restarts the same multi-year qualification cycle.*

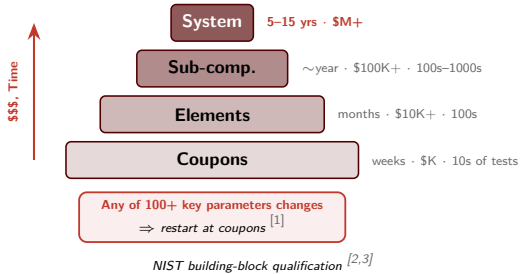

NIST building-block qualification [2,3]

Market is growing fast

Aerospace & Defense AM

$6.1 B (2026) → $11.5 B (2030)

(∼17% CAGR) [4]

Lockheed Martin, EOS, Sintavia all betting on faster process readiness [5].

Why LPBF Qualification Hits a Wall

*Metal AM markets are growing fast — but every **new alloy** and every **parameter change** restarts the same multi-year qualification cycle.*

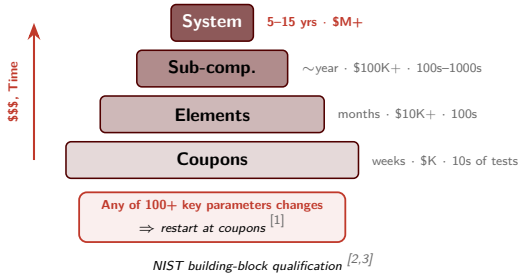

NIST building-block qualification [2,3]

Market is growing fast

Aerospace & Defense AM

$6.1 B (2026) → $11.5 B (2030)

(∼17% CAGR) [4]

Lockheed Martin, EOS, Sintavia all betting on faster process readiness [5].

...but every alloy pays this cost

- Single complex aerospace LPBF component: **$4 M+**, **2+ years** just to qualify [6]
- Standards demand **>99% relative density** [7,8]
- Statistical qualification: **thousands of tests** [3]

Why LPBF Qualification Hits a Wall

*Metal AM markets are growing fast — but every **new alloy** and every **parameter change** restarts the same multi-year qualification cycle.*

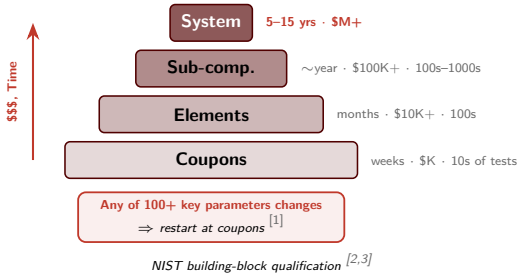

NIST building-block qualification [2,3]

Market is growing fast

Aerospace & Defense AM

$6.1 B (2026) → $11.5 B (2030)

(∼17% CAGR) [4]

Lockheed Martin, EOS, Sintavia all betting on faster process readiness [5].

...but every alloy pays this cost

- Single complex aerospace LPBF component: **$4 M+**, **2+ years** just to qualify [6]
- Standards demand **>99% relative density** [7,8]
- Statistical qualification: **thousands of tests** [3]

Binding constraint is qualification cost — not part complexity, not machine capacity.

The Bet: Physics Can Reduce the Experimental Burden

*Thermal physics of LPBF is **already known** before any build. What if cross-material screening lived in that physics, not in cross-material data?*

From 200-Hour Sweeps to 1.25-Second Predictions

NIST-style parameter sweep for **one** new alloy: $\sim$**200 experimental hours** to map a usable process window.

Test alloys (**Ti6Al4V, CuCrZr**) *not in training*;
they bracket the **widest thermal-diffusivity span** across the six alloys.

The Bet: Physics Can Reduce the Experimental Burden

*Thermal physics of LPBF is **already known** before any build. What if cross-material screening lived in that physics, not in cross-material data?*

From 200-Hour Sweeps to 1.25-Second Predictions

NIST-style parameter sweep for **one** new alloy: $\sim$**200 experimental hours** to map a usable process window.

Test alloys (**Ti6Al4V, CuCrZr**) *not in training*;
they bracket the **widest thermal-diffusivity span** across the six alloys.

Where the "learning" actually lives

Conventional ML for AM:
knowledge $\leftarrow$ $N = 10^3$–10^4 **experiments** per alloy

PIKNN:
knowledge $\leftarrow$ **physics handbook** + K=1–10 **shots/class**

Physics is paid for once, re-used for every new alloy.

The Bet: Physics Can Reduce the Experimental Burden

*Thermal physics of LPBF is **already known** before any build. What if cross-material screening lived in that physics, not in cross-material data?*

From 200-Hour Sweeps to 1.25-Second Predictions

NIST-style parameter sweep for **one** new alloy: ∼**200 experimental hours** to map a usable process window.

Test alloys (**Ti6Al4V, CuCrZr**) *not in training*;
they bracket the **widest thermal-diffusivity span** across the six alloys.

Where the "learning" actually lives

Conventional ML for AM:
knowledge ← $N = 10^3$–10^4 **experiments** per alloy

PIKNN:
knowledge ← **physics handbook** + K=1–10 **shots/class**

Physics is paid for once, re-used for every new alloy.

What this enables

- **Training-free**: fixed physics map + K-shot target-alloy support
- Source alloys: 4 (1,244 builds); held-out: Ti6Al4V, CuCrZr
- **No retraining** when a new alloy arrives
- **Early-stage screening** before certification builds

The Bet: Physics Can Reduce the Experimental Burden

*Thermal physics of LPBF is **already known** before any build. What if cross-material screening lived in that physics, not in cross-material data?*

From 200-Hour Sweeps to 1.25-Second Predictions

NIST-style parameter sweep for **one** new alloy: ∼**200 experimental hours** to map a usable process window.

Test alloys (**Ti6Al4V, CuCrZr**) *not in training*;
they bracket the **widest thermal-diffusivity span** across the six alloys.

Where the "learning" actually lives

Conventional ML for AM:
knowledge ← $N = 10^3$–10^4 **experiments** per alloy

PIKNN:
knowledge ← **physics handbook** + K=1–10 **shots/class**

Physics is paid for once, re-used for every new alloy.

What this enables

- **Training-free**: fixed physics map + K-shot target-alloy support
- Source alloys: 4 (1,244 builds); held-out: Ti6Al4V, CuCrZr
- **No retraining** when a new alloy arrives
- **Early-stage screening** before certification builds

Hypothesis: thermal physics is a **cross-material bridge** — a non-parametric classifier in that space can match or exceed learned embeddings.

The Cross-Material Stress Test

*We hold out the two alloys at the **thermal extremes** of the dataset — a difficult held-out-alloy stress test.*

Dataset — 1,579 builds [9]

Split	Material	Samples
Train (4 alloys)	316L	358
	AlSi10Mg	353
	18Ni300	292
	IN718	241
Test (held-out)	Ti6Al4V	197
	CuCrZr	138

Thermal extremes (m^2/s):
Ti6Al4V 3.0×10^{-6} $\leftrightarrow$ CuCrZr 1.1×10^{-4}
$\Rightarrow$ $\sim$**37**$\times$ contrast at the held-out extremes

The Cross-Material Stress Test

*We hold out the two alloys at the **thermal extremes** of the dataset — a difficult held-out-alloy stress test.*

Dataset — 1,579 builds [9]

Split	Material	Samples
Train (4 alloys)	316L	358
	AlSi10Mg	353
	18Ni300	292
	IN718	241
Test (held-out)	Ti6Al4V	197
	CuCrZr	138

Thermal extremes (m^2/s):
Ti6Al4V 3.0×10^{-6} $\leftrightarrow$ CuCrZr 1.1×10^{-4}
$\Rightarrow$ $\sim$**37**$\times$ contrast at the held-out extremes

Test-Alloy Class Imbalance

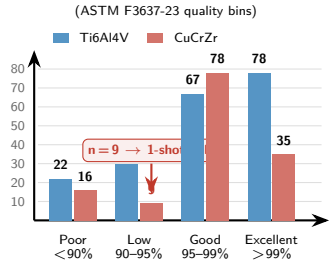

The Cross-Material Stress Test

*We hold out the two alloys at the **thermal extremes** of the dataset — a difficult held-out-alloy stress test.*

Dataset — 1,579 builds [9]

Split	Material	Samples
Train (4 alloys)	316L	358
	AlSi10Mg	353
	18Ni300	292
	IN718	241
Test (held-out)	Ti6Al4V	197
	CuCrZr	138

Thermal extremes (m^2/s):
Ti6Al4V 3.0×10^{-6} $\leftrightarrow$ CuCrZr 1.1×10^{-4}
$\Rightarrow$ $\sim$**37**$\times$ contrast at the held-out extremes

Test-Alloy Class Imbalance

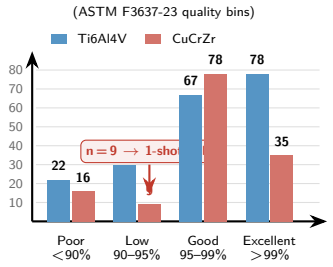

Severe imbalance in low-RD regimes + thermal contrast = **challenging** cross-material benchmark.

Where Existing Methods Fall Short

Existing approaches typically miss at least one requirement for this setting: ***few-shot*** *+* ***explicit physics*** *+* ***no retrain on target*** *+* ***cross-material transfer****.*

Approach	*Few-shot*	*Explicit physics*	*No source meta training retrain*	*AM qual. classif.*	*Cross-m. transfer*
Conventional LPBF ML [10]	✗	✗	✗	✓	✗
FSL / ProtoNet [11,12]	✓	✗	✗	✓	✗
PIML for AM [13,14]	✗	✓	✗	✗	✗
Material-agnostic ML [15]	✗	✗	✗	✗	✓
PIKNN (this work)	✓	✓	✓	✓	✓

The opening

A **physics-informed**, **training-free**, **cross-material** few-shot classifier for LPBF relative density — targeted at **early decision support** under 1–10-shot supervision.

Why all five capabilities matter

A new alloy in qualification has few labels, demands real physics, and cannot afford a model re-train per parameter change.

Approach — Replace Learned Embeddings with Physics

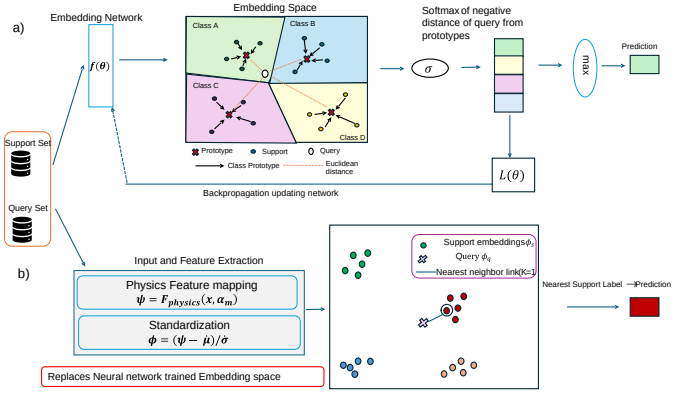

(a) Prototypical Network — baseline

- MLP encoder f_θ: **13,056 parameters**
- $\mathbb{R}^8 \rightarrow$ 64-D **learned** space
- Softmax over prototype distances
- **10,000 meta-training** episodes

(b) PIKNN — proposed

- Operates directly in **8-D physics space**
- **Zero trainable parameters**
- 1-NN: nearest support label, Euclidean distance
- No meta-training, no hyperparameter tuning

Key idea: Replace the **learned embedding** with an **explicit physics-engineered feature space** — both methods share identical preprocessing and episodic protocol.

K = number of labeled support examples per class; PIKNN classifies each query by 1-NN regardless of K.

The 8-D Physics-Informed Feature Space

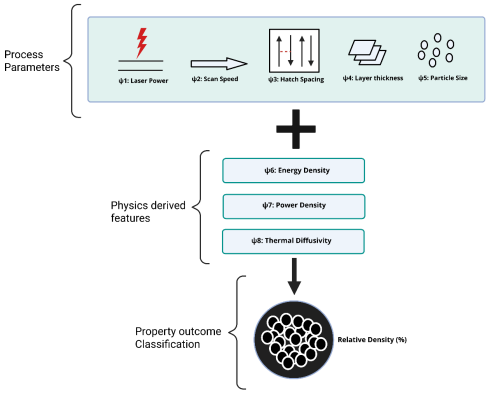

Pipeline: raw process parameters → physics-derived metrics → standardized 8-D vector → RD class prediction.

8-D Feature Vector

$\psi = F_{\text{physics}}(\mathbf{x}, \alpha_m)$

	Feature	Physical role
F1	P (Laser power)	Energy input
F2	v (Scan speed)	Traverse rate
F3	h (Hatch spacing)	Track overlap
F4	l (Layer thickness)	Build resolution
F5	D_{50} (Powder size)	Absorptivity
F6	ED $= \frac{P}{v\,h\,l}$	**Melting completeness**
F7	PD $= \frac{4P}{\pi d^2}$	**Melt-pool regime**
F8	α_m (Thermal diffusivity)	**Cross-material bridge**

Standardize: $\phi = (\psi - \hat{\mu})/\hat{\sigma}$

Predict: $\hat{y} = y_{i^*}$, $i^* = \arg\min_i \|\phi_q - \phi_i\|_2$

$\hat{\mu}$, $\hat{\sigma}$ estimated from source-alloy training data only (no held-out test leakage). α_m for each alloy from literature.

Experimental Setup — Cross-Material Benchmark

Samples per Quality Bin (ASTM F3637-23)

Material	Poor <90%	Low 90–95%	Good 95–99%	Excellent >99%
316L	60	50	161	87
AlSi10Mg	32	65	152	104
18Ni300	13	11	114	154
IN718	7	11	178	45
Ti6Al4V	22	30	67	78
CuCrZr	16	9	78	35

Episodic Protocol

- **4-way** classification (Poor, Low, Good, Excellent)
- Support: $K \in \{1, 2, 3, 5, 10\}$ shots/class
- Query: **8 samples/class/episode**
- **300 episodes** per configuration; 95% bootstrap CIs (1,000 resamples)

Baselines (identical features & protocol)

- Prototypical Network [11]: 13,056 params, 10k meta-training episodes
- SVM (RBF kernel) — classical parametric control

All three methods share preprocessing, features, and episodic structure — isolating **representation quality** as the experimental variable.

Result — PIKNN Beats Trained Baselines

Method	Material	K	Acc. (%)
PIKNN	Ti6Al4V	1	33.1 ± 1.0
	Ti6Al4V	3	42.0 ± 1.2
	Ti6Al4V	5	48.8 ± 1.1
	Ti6Al4V	10	$\mathbf{58.0 \pm 1.0}$
	CuCrZr	1	$\mathbf{52.0 \pm 1.4}$

Result — PIKNN Beats Trained Baselines

Method	Material	K	Acc. (%)
PIKNN	Ti6Al4V	1	33.1 ± 1.0
	Ti6Al4V	3	42.0 ± 1.2
	Ti6Al4V	5	48.8 ± 1.1
	Ti6Al4V	10	$\mathbf{58.0 \pm 1.0}$
	CuCrZr	1	$\mathbf{52.0 \pm 1.4}$
ProtoNet	Ti6Al4V	1	33.6 ± 1.0
	Ti6Al4V	3	36.7 ± 1.1
	Ti6Al4V	5	37.7 ± 1.0
	Ti6Al4V	10	39.7 ± 1.0
	CuCrZr	1	46.1 ± 1.2

Result — PIKNN Beats Trained Baselines

Method	Material	K	Acc. (%)
PIKNN	Ti6Al4V	1	33.1 ± 1.0
	Ti6Al4V	3	42.0 ± 1.2
	Ti6Al4V	5	48.8 ± 1.1
	Ti6Al4V	10	$\mathbf{58.0 \pm 1.0}$
	CuCrZr	1	$\mathbf{52.0 \pm 1.4}$
ProtoNet	Ti6Al4V	1	33.6 ± 1.0
	Ti6Al4V	3	36.7 ± 1.1
	Ti6Al4V	5	37.7 ± 1.0
	Ti6Al4V	10	39.7 ± 1.0
	CuCrZr	1	46.1 ± 1.2
SVM	Ti6Al4V	1	29.1 ± 1.0
	Ti6Al4V	3	35.4 ± 1.2
	Ti6Al4V	5	36.8 ± 1.1
	Ti6Al4V	10	40.3 ± 1.0
	CuCrZr	1	42.0 ± 1.4

Result — PIKNN Beats Trained Baselines

Method	Material	K	Acc. (%)
PIKNN	Ti6Al4V	1	33.1 ± 1.0
	Ti6Al4V	3	42.0 ± 1.2
	Ti6Al4V	5	48.8 ± 1.1
	Ti6Al4V	10	$\mathbf{58.0 \pm 1.0}$
	CuCrZr	1	$\mathbf{52.0 \pm 1.4}$
ProtoNet	Ti6Al4V	1	33.6 ± 1.0
	Ti6Al4V	3	36.7 ± 1.1
	Ti6Al4V	5	37.7 ± 1.0
	Ti6Al4V	10	39.7 ± 1.0
	CuCrZr	1	46.1 ± 1.2
SVM	Ti6Al4V	1	29.1 ± 1.0
	Ti6Al4V	3	35.4 ± 1.2
	Ti6Al4V	5	36.8 ± 1.1
	Ti6Al4V	10	40.3 ± 1.0
	CuCrZr	1	42.0 ± 1.4

Key findings

- **+18.3 pp** over ProtoNet at 10-shot (Ti6Al4V)
- **+5.9 pp** over ProtoNet at 1-shot (CuCrZr)
- Gap **widens** with shots — physics features scale; ProtoNet plateaus ($36.7\% \rightarrow 39.7\%$)
- SVM shows same saturation $\Rightarrow$ **representation, not classifier**
- All methods $\gg$ 25% random chance

Performance is **screening-grade**, not certification-grade.

Why It Works — Representation Preserves Class Structure

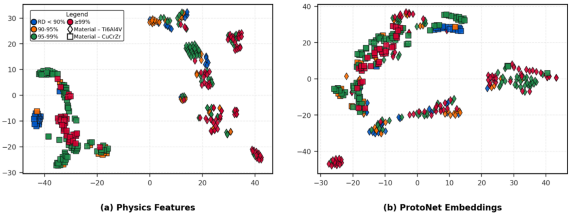

(a) Physics Features (b) ProtoNet Embeddings

t-SNE: physics features (left) form coherent class clusters; ProtoNet embeddings (right) mix classes.

Zero-shot linear probe transfer

Space	Acc. (%)
Physics	45.4
ProtoNet	31.9

+13.5 pp on unseen alloys.

Cluster quality (lower DB, higher CH)

Space	DB↓	CH↑
Physics (Ti6Al4V only)	**4.08**	**8.4**
ProtoNet (Ti6Al4V only)	6.92	5.9
Physics (both materials)	**6.27**	**6.7**
ProtoNet (both materials)	11.65	2.6

41–46% lower DB ⇒ tighter, better-separated clusters.

The physics space **preserves** cross-material discriminative structure; learned embeddings **lose it**.

Which Physics Terms Actually Matter

Energy density is the largest contributor; thermal diffusivity is the cross-material bridge — both are essential.

Ablation (Ti6Al4V, 5-shot)

Feature Set	Acc %	Δ
All 8 physics features	**48.8**	—
w/o Energy Density (F6)	41.2	−7.6
w/o Thermal Diffusivity (F8)	44.1	−4.7
w/o Power Density (F7)	46.2	−2.6
Raw process params only	38.9	−9.9
Physics-derived only (F6–F8)	46.2	−2.6

Which Physics Terms Actually Matter

Energy density is the largest contributor; thermal diffusivity is the cross-material bridge — both are essential.

Ablation (Ti6Al4V, 5-shot)

Feature Set	Acc %	Δ
All 8 physics features	**48.8**	—
w/o Energy Density (F6)	41.2	−7.6
w/o Thermal Diffusivity (F8)	44.1	−4.7
w/o Power Density (F7)	46.2	−2.6
Raw process params only	38.9	−9.9
Physics-derived only (F6–F8)	46.2	−2.6

Distance contribution (pooled)

Feature	% of dist.
Energy Density (ED)	**24.5%**
Thermal Diffusivity (α_m)	**20.8%**
Hatch Spacing (h)	20.3%
Powder Size (D_{50})	9.7%
Laser Power (P)	7.9%
Power Density (PD)	7.7%
Layer Thickness (l)	5.8%
Scan Speed (v)	3.4%

No single feature $> 25\% \Rightarrow$ balanced encoding.

Which Physics Terms Actually Matter

Energy density is the largest contributor; thermal diffusivity is the cross-material bridge — both are essential.

Ablation (Ti6Al4V, 5-shot)

Feature Set	Acc %	Δ
All 8 physics features	**48.8**	—
w/o Energy Density (F6)	41.2	−7.6
w/o Thermal Diffusivity (F8)	44.1	−4.7
w/o Power Density (F7)	46.2	−2.6
Raw process params only	38.9	−9.9
Physics-derived only (F6–F8)	46.2	−2.6

Distance contribution (pooled)

Feature	% of dist.
Energy Density (ED)	**24.5%**
Thermal Diffusivity (α_m)	**20.8%**
Hatch Spacing (h)	20.3%
Powder Size (D_{50})	9.7%
Laser Power (P)	7.9%
Power Density (PD)	7.7%
Layer Thickness (l)	5.8%
Scan Speed (v)	3.4%

No single feature $> 25\% \Rightarrow$ balanced encoding.

Within an episode, α_m is constant and contributes **0%** to distance — classification is driven entirely by **process parameters and energy metrics**.

Cost Profile — What You Pay to Run It

PIKNN eliminates the entire meta-training phase. Per-episode inference is essentially free.

	ProtoNet	PIKNN
Trainable parameters	13,056	0
Meta-training time	11.9 s + 10k ep.	none
Hyperparameters	lr, batch, dim, ...	$K = 1$ only
Memory (per episode)	encoder weights	≤ 40 vectors

*Total wall-clock: pipeline runs end-to-end in **1.25 s** per new-alloy prediction.*

Why this matters operationally

- **No GPU** needed — runs on CPU/edge
- **No retraining** when a new alloy arrives — only literature lookup of α_m
- **No hyperparameter tuning** — removes a re-validation step every deployment
- Inline screening at $<$**100 ms** feasible

The contribution is not "training-free"; it's that **physics features carry the inductive bias**.

What This Enables in Practice

From a six-alloy benchmark to three concrete deployment paths.

1. Rapid new-alloy onboarding

Input: one α_m value from a materials datasheet.

Output: quality-bin prediction from 1–10 labeled builds.

Reduces the number of candidates entering a physical DOE.

2. In-silico parameter screening

Sweep **1,000s of** (P, v, h, l) **candidates** virtually at 1.25 s each.

Pre-rank the candidates before any physical build.

1,000 candidates: ∼21 min screening vs. ∼200 hrs of LPBF time.

3. Inline screening on the factory floor

$<$100 ms prediction $\rightarrow$ feasible **between layers / between builds**.

Runs on **CPU/edge** hardware. No model server.

Quality decisions move from post-build inspection to in-process.

Where this fits in the qualification stack: PIKNN does **not** replace ASTM/AS9100 certification — it **narrows the experimental window** that certification testing must cover.

Takeaways

What we showed

1. Cross-material few-shot benchmark for LPBF RD across 6 alloys, with Ti6Al4V & CuCrZr at the thermal extremes.
2. Physics features preserve class structure: **41–46% lower DB** than trained embeddings.
3. Training-free PIKNN beats a 13k-parameter ProtoNet by **up to +18.3 pp**; gap *widens* with shots.

Limitations & future work

- Works because LPBF physics is well-known — not portable to physics-poor regimes.
- Screening-grade, not certification-grade.
- Future: leave-one-out across all 6 alloys; weighted distance metrics; in-situ melt-pool fusion.

Key message

When LPBF thermophysics is well-characterized, explicit physics encoding can substitute for data-driven representation learning.

Thank you — questions?

pdhakal@tamu.edu

References I

[1] Federal Aviation Administration. Advisory circular AC 33.15-3: Powder bed fusion additive manufacturing process for aircraft engine components. U.S. Department of Transportation, Federal Aviation Administration, 2024. URL https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_33.15-3.pdf. Accessed: 2026-05.

[2] Mohsen Seifi, Ayman Salem, Jack Beuth, Ola Harrysson, and John J. Lewandowski. Overview of materials qualification needs for metal additive manufacturing. *JOM*, 68(3):747–764, 2016. doi: 10.1007/s11837-015-1810-0.

[3] National Institute of Standards and Technology. Qualification for additive manufacturing materials, processes, and parts. NIST Programs and Projects, 2022. URL https://www.nist.gov/programs-projects/qualification-additive-manufacturing-materials-processes-and-parts. Accessed: 2026-05.

[4] Research and Markets. Aerospace and defense additive manufacturing market report 2026: $11.48 Bn opportunities, trends, competitive landscape, strategies, and forecasts, 2020-2025, 2025-2030F, 2035F. GlobeNewswire press release, January 2026. URL https://www.globenewswire.com/news-release/2026/01/20/3222099/28124/en/Aerospace-and-Defense-Additive-Manufacturing-Market-Report-2026-11-48-Bn-Opportunities-Trends-Competitive-Landscape-Strategies-and-Forecasts-2020-2025-2025-2030F-2035F.html. Accessed: 2026-05.

[5] Lockheed Martin Corporation. Accelerating operational readiness with laser powder-bed fusion additive manufacturing. Lockheed Martin News Feature, 2026. URL https://www.lockheedmartin.com/en-us/news/features/2026/Accelerating-Operational-Readiness-with-Laser-Powder-Bed-Fusion-Additive-Manufacturing.html. Accessed: 2026-05.

[6] David Paul Spadaccia. Laser powder bed fusion of a complex aerospace component: Effects of layerwise process control on part quality. Master's thesis, Virginia Polytechnic Institute and State University, June 2025. URL https://vtechworks.lib.vt.edu/items/ed197e3a-0007-4919-8f15-9ea97cf6dc9e. VTechWorks open repository.

[7] ASTM International. F3637-23 standard guide for additive manufacturing of metal — finished part properties — methods for relative density measurement. ASTM International, West Conshohocken, PA, 2023. URL https://store.astm.org/f3637-23.html.

[8] SAE International. AMS7003: Laser powder bed fusion process. SAE Aerospace Material Specification, 2023. URL https://saemobilus.sae.org/standards/ams7003-laser-powder-bed-fusion-process. Accessed: 2026-05.

[9] German O. Barrionuevo, Iván La Fé-Perdomo, and Jorge A. Ramos-Grez. Laser powder bed fusion dataset for relative density prediction of commercial metallic alloys. *Scientific Data*, 12: 375, 2025. doi: 10.1038/s41597-025-04576-x.

[10] Chengcheng Wang, X. P. Tan, S. B. Tor, and C. S. Lim. Machine learning in additive manufacturing: State-of-the-art and perspectives. *Additive Manufacturing*, 36:101538, 2020. doi: 10.1016/j.addma.2020.101538.

[11] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning (ICML)*, 2017.

[13] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021. doi: 10.1038/s42254-021-00314-5.

[14] Sina Faegh et al. Physics-informed machine learning for additive manufacturing process modeling and quality prediction. *Journal of Manufacturing Processes*, 2025. In press.

[15] Anqi Wang, Heran Song, and Qiang Huang. Material-agnostic machine learning for laser powder bed fusion process modeling. *Journal of Manufacturing Systems*, 2023.