

Which Industry 4.0 investment comes first?

A decision-support framework for sequencing Industry 4.0 investments in lean surface-mount electronics production

● iPhone main-logic-board assembly

● China & India high-volume EMS

MOTIVATION

A plant manager has £100,000 and four good options.

IoT-enabled SPC. MES. Machine-learning AOI. A digital twin. Every vendor promises a return. Every case study is a success story.

THE LITERATURE SAYS

Industry 4.0 and Lean reinforce each other — nineteen sources agree.


THE LITERATURE DOES NOT SAY

Which one to buy *first*, on *this* line, with *this* budget.

Buy the wrong thing first and you don't get a smaller return. You very often get none at all.

Fleet-scale stakes

India's share of global iPhone output and Tata Electronics' share of India's exports



Sources as cited in §1.1; editable draw.io source: Fig1.2_India_iPhone_Rampup.drawio

INDIA SHARE OF GLOBAL OUTPUT


25%

≈55 m units in 2025, +53% year on year

TATA SHARE OF INDIA EXPORTS


≈37–40%

from 13% in 2024




Tens to hundreds of SMT lines per campus, across two countries, renewed with every new product introduction. The sequencing decision is made hundreds of times a year.

Fourteen stations. Three inspection gates.



Two parallel sub-lines converge where the sub-boards are joined. Solder defects — bridges, insufficient paste, cold joints, voids — dominate the defect Pareto, which is why the inspection-intensive steps carry disproportionate weight and why inspection technology sits at the centre of this project.

A protocol fixed before searching



26
AT FULL TEXT

19
RETAINED

16
PEER-REVIEWED

THE GAP, STATED PRECISELY

The literature establishes **that** Lean and I4.0 reinforce one another, and **which** technologies help which KPIs. It does not establish **in what order** a capital-constrained line should adopt them — conditional on the maturity that line has actually reached.

In what order should IoT-SPC, MES, ML-AOI and a digital twin be introduced on a high-volume smartphone SMT line — given its Lean maturity and its readiness — to get the most return per pound?

SQ1 · MATURITY

What minimum Lean maturity does each technology need before it can pay off?

SQ2 · READINESS

Which organisational, technical or financial factors are the binding bottlenecks?

SQ3 · ROBUSTNESS

How sensitive is the order to capital cost and to assumed KPI improvement?

Five locked objectives, O1–O5, each with its verification evidence fixed in advance.

Design science, pre-registered

Build an artefact, then evaluate it against criteria fixed *before* the evaluation runs.

1

Foundation

PRISMA-informed review; KPI baselines; the gap evidenced.

2

Framework

Maturity rubric, readiness rubric, technology-to-KPI mapping. Supervisor gate.

3

Tool

Rule engine and ROI model, verified against five hand-calculated profiles.

4

Validation

Blind case run, sensitivity analysis, three independent cross-checks.



VERIFICATION BEFORE CONSTRUCTION

Five synthetic profiles were hand-derived from the rule set *before* any code was written. The engine was accepted only by matching an answer fixed in advance.

ACCEPTANCE CRITERIA PRE-REGISTERED

Case-admissibility rules, agreement metrics and sensitivity ranges were published before execution — which forecloses post-hoc tuning.

The Digital-Lean Synergy Model



GATE

Maturity decides what is *eligible*

MODULATE

Readiness decides *when* — deploy, stage or defer

RANK

Return per pound of capital decides the *order*

Maturity is a gate, not a modifier

TECHNOLOGY	MINIMUM LEVEL	WHY
IoT-enabled SPC	L2 Managed	Needs a stable, charted process to control
MES	L2 Managed	Needs standard work to execute against
ML-based AOI	L3 Defined	Needs an adjudicated defect dataset to learn from
Digital twin	L3 Defined	Needs trustworthy process data to mirror

General	Definition	Format	User properties	Value
Expected utilities for different policies:				
Lean maturity (..	L1	L2	L3	
IoT_SPC	70585	70585	70585	
ML_AOI	-140000	-140000	72284	
MES	-200000	12284	12284	
Digital_Twin	-140000	-140000	-17937	

Expected utility by maturity state — solved influence diagram

READ THE MIDDLE ROW

IoT-SPC returns **£70,585 at every level**. ML-AOI returns – **£140,000** at L1 and L2 — the full capital, stranded — and only turns positive at L3.

Below its gate a technology delivers **no** benefit, not a smaller one. Rule R1: at L1 Reactive, nothing is eligible.

Readiness sets the timing

ACTION

STAGE

MATURITY

L2

LIMITING DIMENSION

Technical

R_min = 2

COMPOSITE READINESS

3.00

R_mean (secondary)

L2 eligible set is IoT-SPC, MES; readiness 3.00 with Technical limiting (min 2) → stage; weakest KPI 'ppm' maps to IoT-SPC first (R5); ranked by NPV per pound of capex (R6). Remediate Technical readiness before the next step.

Recommended sequence

Step	Technology	Capex (£)	Net annual (£)	Payback (mo)	NPV 5-yr (£)	IRR	NPV per £	KPI match	Within budget
1	IoT-SPC	100,000	45,000	26.7	70,585	34.9%	0.71	yes	yes
2	MES	200,000	56,000	42.9	12,284	12.4%	0.06	—	yes

DEFER

min = 1 or mean < 2.5

STAGE

remediate first


DEPLOY

min ≥ 3, mean ≥ 3.5

THE BOTTLENECK, MADE VISIBLE

Composite readiness 3.00 looks like a green light. But **Technical limits at 2** — so the rule returns STAGE and names the remediation before capital is released. Not "yes". Not "no". "Fix this first."

From a described framework to an inspectable one




Python engine · Excel ledger · Streamlit interface · MATLAB analytics · GeNle model. The deployed instance **imports the same engine module** as every reported result — so any divergence between the report and the live tool would be immediately visible.

[dlsm-dashboard...streamlit.app](#)

Digital–Lean Synergy Decision-Support Tool


Three coordinated layers, one shared logic — checked by three independent methods

THREE COORDINATED LAYERS · ONE SHARED LOGIC



• **Verification-first.** Five synthetic line profiles were hand-derived from the rule set before coding — **100% action agreement** (Table 25). The one divergence, Synthetic-B, resolved in the tool's favour and the hand note was corrected. Python engine ≡ Excel ledger, exactly.

INDEPENDENT CROSS-VALIDATION → REPORTED IN §4.6




CONVERGENT RECOMMENDATION · VERIFIED INPUTS ONLY

IoT-SPC > ML-AOI > MES > Digital twin


RESULTS

The recommended sequence

Payback, months (base priors)



Five-year NPV at 10% discount




#	TECHNOLOGY	PAYBACK	5-YR NPV	NPV / £
1	IoT-SPC	26.7 mo	£70,585	0.71
2	ML-AOI	30.0 mo	£72,284	0.52
3	MES	42.9 mo	£12,284	0.06
4	Digital twin	52.2 mo	-£17,937	-0.06

ML-AOI has the higher *absolute* NPV. IoT-SPC still ranks first — it takes far less capital to get there, and it is eligible at every maturity level.

Payback (months) and five-year NPV at 10% discount

ANSWERING SQ3


The ranking survives the wobble



1,000 draws · capex $\pm 30\%$ · KPI uplift $\pm 20\%$ · discount 8–12%

STABILITY

IoT-SPC stays fastest and the digital twin slowest across the whole sampled space. The framework outputs a *sequence*, and the sequence is what survives.



One-at-a-time screen for IoT-enabled SPC


WHICH NUMBER TO ARGUE ABOUT

Benefit fraction and capex dominate; the discount rate barely registers. So challenge the benefit fraction first in a capital review.

CROSS-VALIDATION

Four methods, four logics, one answer

RULE ENGINE



Ranks by NPV per pound of capital

ENTROPY-WEIGHTED MCDA

```
e@Mac ~ % python3 mcda_ahp_topsis.py
rary/Developer/CommandLineTools/usr/bin/python3: can't open f
e@Mac ~ % cd "/Users/shree/Downloads/Dessertation/Dissertation
e@Mac Python Tool % ls
_Influence_Diagram.xdsl  README.md
_ahp_topsis.py          smt_line_des.py
e@Mac Python Tool % python3 mcda_ahp_topsis.py

=====
WEIGHTING & RANKING OF I4.0 TECHNOLOGIES (EG7302, verified)
=====

Decision matrix (Core File Table 4.5):
Technology      Payback (months)  NPV 5-yr (GBP)
IoT-enabled SPC      26.7              70,585
ML-based AOI        30.0              72,284
MES                 42.9              12,284
Digital twin        52.2             -17,937

Entropy-derived criterion weights (from the data above):
Payback (months)    0.390
NPV 5-yr (GBP)      0.293
IRR (%)             0.317

=====
Ranking results:
=====
Technology      TOPSIS rank  VIKOR Q rank  PROM.II rank
IoT-enabled SPC    0.991      1      0.000      1      0.804      1
ML-based AOI      0.860      2      0.135      2      0.529      2
MES               0.326      3      0.654      3     -0.333      3
Digital twin      0.000      4      1.000      4     -1.000      4

=====
Sensitivity analysis:
=====
Sensitivity sequence (mean rank across TOPSIS / VIKOR / PROMETHEE II)
. IoT-enabled SPC (mean rank 1.00)
. ML-based AOI (mean rank 2.00)
. MES (mean rank 3.00)
. Digital twin (mean rank 4.00)

=====
Recommended order (tool.py, for comparison): IoT-SPC -> ML-AOI -> MES -> Digital twin
=====
```

TOPSIS · VIKOR · PROMETHEE II — weights derived from the data

INFLUENCE DIAGRAM



Expected utility, maturity treated as uncertain

DISCRETE-EVENT SIMULATION

```
T FRONT-END LINE - DISCRETE-EVENT SIMULATION (SimPy) EG7302
RT 1 of 2 — single seeded run (as reported in Core File Table 4.15)

=====
Configuration Summary
=====
Configuration      FPY %  Escaped DPPM  FalseCall PPM  Boards/hr
Baseline           96.09      1578          28409          396.0
+ IoT-SPC          97.41         631          28093          396.0 (FPY +1.33)
+ IoT-SPC + ML-AOI 97.41         0           7891          396.0 (FPY +1.33)

=====
RT 2 of 2 — verification & validation protocol (Sec. 3.6.6)
=====
VERIFICATION CHECKS (Sargent, 2013; Law, 2015) — degenerate cases
=====
V1 zero defect rate -> FPY=100.00% escapes= 0 DPPM PASS (expect FPY=100.0)
V2 perfect inspection -> FPY= 96.73% escapes= 0 DPPM PASS (expect escapes=0)

=====
Verification result: ALL CHECKS PASSED
=====
-REPLICATION ANALYSIS (warm-up = 30 min excluded; 95% CI, Student's t, df=29)
=====
Warm-up was determined by inspection of the cumulative-output series (Law, 2015): mean cycles. A 30-minute warm-up (>> 4.0 min) is a conservative exclusion window relative to the mean cycle time.

=====
Configuration Summary (Mean and 95% CI)
=====
Configuration      FPY % (mean±CI)  Escaped DPPM  Boards/hr
Baseline           95.83 +/-0.15    967 +/-180     399.6 +/-0.4
+ IoT-SPC          97.25 +/-0.10    667 +/-131     399.6 +/-0.4
+ IoT-SPC + ML-AOI 97.25 +/-0.10    233 +/-114     399.6 +/-0.4

- -> S2 mean FPY gap: +1.43 pts; 95% CIs DO NOT OVERLAP (effect is statistically distinct)

=====
Definitions:
=====
Y = boards built defect-free (true first-pass yield).
Escaped DPPM = defective boards passing post-AOI undetected (outgoing quality).
1 INPUTS are declared parameters/design priors (see header);
1 OUTPUTS above are COMPUTED by the model, not hand-entered.

=====
References (Harvard, University of Leicester convention)
=====
Banks, J., Carson, J.S., Nelson, B.L. and Nicol, D.M. (2010) Discrete-Event System Simulation. 5th edn. Upper Saddle River, NJ: Pearson.
Law, A.M. (2015) Simulation Modeling and Analysis. 5th edn. New York: McGraw-Hill.
Matloff, N. (2008) Introduction to Discrete-Event Simulation and the SimPy Language. Davis, CA: University of California, Davis.
Sargent, R.G. (2013) 'Verification and validation of simulation models', Journal of Simulation, 7(1), pp. 12-24.
Team SimPy (2023) SimPy: Discrete event simulation for Python (Version 4). Available at: https://simpy.readthedocs.io.
```

30 seeded replications of the SMT front end

IoT-SPC > ML-AOI > MES > digital twin — reproduced independently by all four.

What the evidence does — and does not — support

ESTABLISHED

Engine verified 5/5. Ledger agrees exactly. Ranking stable across 1,000 draws. Four independent methods converge.

CORROBORATED, ONCE

Four published cases screened against pre-registered criteria — **one** admissible. The blind run confirmed the gate and exposed a rule defect, since corrected.

NOT ESTABLISHED

Field performance. Inter-rater reliability of the rubrics — protocol pre-registered, deferred pending ethical approval rather than rushed.

```
shree@Mac ~ % python3 mcda_ahp_topsis.py
/Library/Developer/CommandLineTools/usr/bin/python3: can't open file '/Users/shree/mcda_ahp_topsis.py':
shree@Mac ~ % cd "/Users/shree/Downloads/Dessertation/Dissertation Draft/Python Tool"
shree@Mac Python Tool % ls
DLSM_Influence_Diagram.xdsl  README.md
mcda_ahp_topsis.py           smt_line_des.py
shree@Mac Python Tool % python3 mcda_ahp_topsis.py

[MCDA WEIGHTING & RANKING OF I4.0 TECHNOLOGIES (EG7302, verified data)]
=====
Decision matrix (Core File Table 4.5):
[Technology      Payback (months)  NPV 5-yr (GBP)  IRR (%)
IoT-enabled SPC    26.7             70,585         34.9
ML-based AOI      30.0             72,284         28.6
MES               42.9             12,284         12.4
Digital twin      52.2            -17,937          4.8
=====
Entropy-derived criterion weights (from the data above):
Payback (months)    0.390
NPV 5-yr (GBP)     0.293
IRR (%)             0.317
=====
Technology      TOPSIS  rank  VIKOR Q  rank  PROM.II  rank
=====
IoT-enabled SPC    0.991    1    0.000    1    0.804    1
ML-based AOI      0.860    2    0.135    2    0.529    2
MES               0.326    3    0.654    3   -0.333    3
Digital twin      0.000    4    1.000    4   -1.000    4
=====
Consensus sequence (mean rank across TOPSIS / VIKOR / PROMETHEE II):
1. IoT-enabled SPC (mean rank 1.00)
2. ML-based AOI   (mean rank 2.00)
3. MES            (mean rank 3.00)
4. Digital twin   (mean rank 4.00)

Payback-only order (tool.py, for comparison): IoT-SPC -> ML-AOI -> MES -> Digital twin
```

THE RESULT I'M MOST GLAD I KEPT

A single seeded run showed **zero** escaped defects under S2. Thirty replications showed that was one lucky draw — the real figure is a **75.9% reduction**, not elimination. I report both. The discrepancy is evidence the V&V protocol worked: it caught an overclaim before it reached the discussion.

Four methods agreeing is a statement about *coherence* — and coherence is compatible with being consistently wrong. This is corroboration, not validation.

What is new here

Sequencing, not selection

The literature ranks technologies by benefit. This framework ranks them by *order of adoption*, conditional on the maturity a line has actually reached.

A gate, not a weighting

Maturity is a hard eligibility condition. Below the gate the benefit is not discounted — it is absent, and the capital is stranded.

An instruction, not a verdict

The output names the limiting dimension and the remediation required, so a "not yet" arrives with the reason and the fix attached.

An artefact you can run

Deployed, browser-reachable, importing the verified engine. The framework is inspectable rather than merely described.

FUTURE WORK

What I would do next, in order

FW1 · FIRST

Run the reliability study

Three raters, eight vignettes.
Designed, pre-registered, script written — it needs approval and raters, nothing more.

FW2

Replace priors with pilot data

The benefit fraction is the parameter the sensitivity analysis says matters most. One instrumented line would settle it.

FW3

Widen the case base

One admissible case cannot carry external validity. Industry access would allow a genuine multiple-case design.

FW4 · FW5

Extend and generalise

Multi-line capital allocation across a campus; and test whether the gate logic transfers beyond SMT.

Each direction is tied to a specific limitation, with the method that would address it named.

CONCLUSION

Stabilise, then sense, then automate.

For a line at or above L2 Managed, with defect PPM as its weakest KPI and a constrained budget:

1 · IoT-enabled SPC



2 · ML-based AOI



3 · MES



4 · Digital twin

Robust to the choice of appraisal metric, to the aggregation rule, to uncertainty about maturity, and to a thousand-draw Monte Carlo.

Tested against one real deployment — which corroborated the gate and found a defect. That qualification travels with the answer.

REFERENCES

Sources cited in this presentation

Bloomberg (2026) *Apple now makes about 25% of iPhones in India after China pivot.*

<https://www.bloomberg.com>

Buer, S.-V., Strandhagen, J.O. and Chan, F.T.S. (2018) 'The link between Industry 4.0 and lean manufacturing', *International Journal of Production Research*, 56(8), pp. 2924–2940.

<https://doi.org/10.1080/00207543.2018.1442945>

Free Press Journal (2026) *Tata Electronics overtakes Foxconn in iPhone exports under India's PLI scheme.* <https://www.freepressjournal.in>

Gregor, S. and Hevner, A.R. (2013) 'Positioning and presenting design science research for maximum impact', *MIS Quarterly*, 37(2), pp. 337–355.

Hevner, A.R. et al. (2004) 'Design science in information systems research', *MIS Quarterly*, 28(1), pp. 75–105. <https://doi.org/10.2307/25148625>

Howard, R.A. and Matheson, J.E. (2005) 'Influence diagrams', *Decision Analysis*, 2(3), pp. 127–143. <https://doi.org/10.1287/deca.1050.0020>

Hwang, C.-L. and Yoon, K. (1981) *Multiple attribute decision making*. Berlin: Springer. <https://doi.org/10.1007/978-3-642-48318-9>

IPC (2022) *Quality benchmarks for the electronics assembly industry*. Bannockburn, IL: IPC International.

Lotfi, F.H. and Fallahnejad, R. (2010) 'Imprecise Shannon's entropy and multi attribute decision making', *Entropy*, 12(1), pp. 53–62. <https://doi.org/10.3390/e12010053>

MeitY (2025) *Production-linked incentive (PLI) scheme for large-scale electronics manufacturing*. New Delhi: Government of India.

Opricovic, S. and Tzeng, G.-H. (2004) 'Compromise solution by MCDM methods', *European Journal of Operational Research*, 156(2), pp. 445–455. [https://doi.org/10.1016/S0377-2217\(03\)00020-1](https://doi.org/10.1016/S0377-2217(03)00020-1)

Page, M.J. et al. (2021) 'The PRISMA 2020 statement', *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>

Peffer, K. et al. (2007) 'A design science research methodology for information systems research', *Journal of Management Information Systems*, 24(3), pp. 45–77. <https://doi.org/10.2753/MIS0742-1222240302>

Saltelli, A. et al. (2008) *Global sensitivity analysis: the primer*. Chichester: Wiley. <https://doi.org/10.1002/9780470725184>


Sargent, R.G. (2013) 'Verification and validation of simulation models', *Journal of Simulation*, 7(1), pp. 12–24. <https://doi.org/10.1057/jos.2012.20>

Schumacher, A., Erol, S. and Sihn, W. (2016) 'A maturity model for assessing Industry 4.0 readiness', *Procedia CIRP*, 52, pp. 161–166. <https://doi.org/10.1016/j.procir.2016.07.040>

Team SimPy (2023) *SimPy: discrete-event simulation for Python (Version 4)*. <https://simpy.readthedocs.io>

Tonelli, F. et al. (2022) '3D-CUBE readiness model for Industry 4.0', *Production & Manufacturing Research*, 10(1), pp. 989–1018. <https://doi.org/10.1080/21693277.2022.2135628>

Venable, J., Pries-Heje, J. and Baskerville, R. (2016) 'FEDS: a framework for evaluation in design science research', *European Journal of Information Systems*, 25(1), pp. 77–89. <https://doi.org/10.1057/ejis.2014.36>



Thank you

Questions

Live tool · [dlsn-dashboard...streamlit.app](https://dlsn-dashboard.streamlit.app)

Code & data · [GitHub repository](#)