Soft Computing Fusion with Applications

www.scfa.reapress.com

Soft. Comput. Fusion. Appl. Vol. 3, No. 3 (2026) 208–220.

Paper Type: Original Article

Preview Based View Selection and Cost Aware Stopping for Budgeted Inspection

Shaban Kaji[1,*], Min-Chul Park[2]

[1] Department of Civil Engineering, North Carolina Agricultural and Technical State University, Greensboro, North Carolina, USA; kaji@ncat.edu.

[2] Department of Mechanical Design Engineering, Korea Advanced Institute of Engineering Design (KAIED), Daejeon, Republic of Korea; mcpark@kaied.ac.kr.

Citation:

Received: 06 October 2025 *Revised: 25 February 2026* *Accepted: 21 April 2026*	Kaji, Sh., & Park, M. (2026). Preview based view selection and cost aware stopping for budgeted inspection. *Soft Computing Fusion with Applications, 3*(3), 208-220.

Abstract

Budgeted inspection combines evidence acquisition with a decision about when to stop. This manuscript proposes AVID, a partially observed decision architecture, and evaluates a simpler committee-based realization on 1,797 handwritten-digit images. Classes {3,8,9} represent a proxy defect label, and six deterministic image transforms represent views. Across twenty repeated splits, the preview-based selector reports single-view AUC 0.952 versus 0.864 for random order. Its candidate scoring uses actual unacquired view features, giving it information unavailable to the non-preview baselines. Cost-aware stopping uses 1.65 acquired views per decision versus 3.90 for a posterior-odds threshold baseline, but has lower reported net utility: −0.058 versus +0.036 under correct=+1, wrong=−5, and acquisition cost 0.15. Always acquiring six views has accuracy 0.9392 versus 0.8650 with stopping. The result therefore documents a preview-assisted selection and accuracy–cost trade-off on a digit proxy, not validated industrial acquisition efficiency. Reinforcement learning, attention-based aggregation, asymmetric error costs, and formal Wald guarantees are not executed by this realization.

Keywords: Active vision, Industrial inspection, Anomaly detection, Epistemic uncertainty, Reinforcement learning, Optimal stopping.

1 | Introduction

On a printed-circuit-board assembly line, an automated optical inspection station has a few hundred milliseconds per board. Each additional camera angle, lighting change, or zoom step consumes cycle time, and the station's throughput bounds the throughput of the line; yet a single fixed view is often insufficient, since hairline cracks, casting shrinkage on side walls, and contaminants under components reveal themselves only from particular viewing and illumination geometries. Human inspectors behave differently: they tilt the

Corresponding Author: kaji@ncat.edu

https://doi.org/10.22105/scfa.v3i3.111

Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0).

board, change the angle of the light, look again at ambiguous regions, and stop when the evidence is sufficient. That adaptive, evidence-driven acquisition behaviour, choosing where to look next and when to stop, is precisely what automated inspection lacks. The economic stakes sit in the asymmetry between the two errors: a false accept ships a defective part and incurs field-failure and recall costs, while a false reject scraps good material and interrupts supply; both errors should be priced explicitly, and acquisition itself has a price.

The component literatures are mature. Active perception, articulated by Bajcsy and formalized in the active vision programme of Aloimonos and colleagues, established that sensing should be treated as a controlled process rather than a fixed channel [1], [2]. Industrial anomaly detection now has strong public benchmarks, most prominently the MVTec Anomaly Detection dataset, and a rich method repertoire spanning reconstruction-based, distribution-based, and memory-based models [3–7]. Uncertainty estimation offers calibrated proxies for what a model does not know, with deep ensembles the practical standard [8–10]. Reinforcement learning contributes both the glimpse-policy tradition, in which recurrent agents learn where to look [11], [12], and mature optimizers for sequential decisions [13], [14]. Sequential hypothesis testing contributes the classical machinery of optimal stopping [15], and Bayesian experimental design the expected-information principle [16].

The design question is how to coordinate view selection, evidence accumulation, stopping, and decision costs under a finite budget. Existing approaches address different parts of that question. The present experiments evaluate simple policies on a digit proxy and expose an information-access asymmetry: preview-based candidates are evaluated using features before the acquisition cost is charged. A deployment interpretation therefore requires a priced or shared preview mechanism.

AVID is proposed as a budgeted partially observed architecture with acquisitions and accept, reject, or escalate decisions. The full design includes an ensemble scorer, evidence aggregator, and learned policy. Section 5 uses a supervised logistic-regression committee, deterministic transformed views, and handcrafted selection/stopping rules. The neural architecture and reinforcement-learning policy remain design proposals.

The formulation prices acquisitions and terminal decisions within a finite budget (Section 3). The proposed architecture combines evidence scoring, aggregation, and joint acquisition/stopping actions; a sequential-ratio heuristic provides an audit comparator (Section 4). The formal information-gain objective integrates over prospective observations. The executed preview score is a different, directly evaluated entropy heuristic (Sections 3.3 and 5.1). The reported comparison uses twenty repeated splits for each of three proxy defect-class compositions. Negative utility results and preview access define the empirical scope (Section 5). Assumptions and failure conditions distinguish the proposed industrial architecture from the digit-based realization.

The manuscript is organized as follows. Section 2 synthesizes related work across active vision, industrial anomaly detection, and sequential decision-making. Section 3 states the research questions, hypotheses, and formal problem. Section 4 develops the architecture, training procedure, and Algorithm 1. Section 5 reports the executed experimental study and its analysis, and Sections 6 and 7 close with limitations, counterarguments, and the roadmap.

2 | Related Work

2.1 | Active Vision, Attention, and Experimental Design

Active perception treats sensing as an action rather than a fixed input [1], [2]. Learned glimpse policies [11], [12] and attention mechanisms [17] offer possible architectures, but their learning objectives and sensing assumptions differ. In the present design, acquisition and terminal decisions share a budget; the executed proxy study uses committee scores and one-step rules rather than a trained glimpse or attention policy.

Bayesian experimental design values an acquisition by integrating posterior information over possible observations [16]. Active learning uses related criteria for label queries [18]. Some acquisition methods are

myopic and others plan over budgets; the entire literature cannot be characterized as ignoring future opportunities. The proposed learned policy would address sequential value, while the executed study uses one-step rules.

2.2 | Industrial Anomaly Detection

Industrial defect detection was transformed by the availability of realistic public data, above all the MVTec Anomaly Detection dataset with its fifteen object and texture categories and pixel-precise ground-truth masks [3], later extended in a journal treatment [19] and to logical anomalies [20]. Method families are complementary: GANomaly trains a generative model to reconstruct only nominal appearances and detects deviations in latent space [4]; PaDiM models per-patch distributions of pretrained features [5]; and PatchCore retrieves from a memory of nominal patches, scoring test patches by core-set distance [6]. A unifying review situates these methods against classical shallow baselines and catalogues their assumptions [7]. Backbone capacity matters, with residual and very deep convolutional networks the default feature extractors [21], [22], and robustness of such pipelines to common corruptions is an active concern [23].

Industrial inspection often uses a station-defined acquisition arrangement; the cited anomaly detectors evaluate images after acquisition [3–7]. Corruption studies [23] can assess robustness to input changes, whereas the present decision problem also chooses whether to acquire another view. In Section 5, handwritten-digit images act as proxy parts and deterministic crops or downsamples as distinct information channels. This demonstrates a comparison within that proxy, not transfer to an industrial imaging station.

2.3 | Sequential Decisions, Stopping, and Uncertainty

Sequential acquisition can be formulated through a partially observed process [24]. Reinforcement-learning algorithms [13], [14], [25] and exploration criteria [26–28] motivate possible policies, but none is trained in this study. Wald's test uses accumulated likelihood ratios under specified sequential sampling assumptions [15]. A discriminative posterior computed from overlapping views is not automatically that likelihood-ratio process. Cost-sensitive learning supplies decision-cost modeling [29]; the executed payoff prices the two error types equally.

Ensembles [8], dropout [9], and calibration methods [10] offer practical evidence models. Member disagreement is a proxy and can be small when every member is wrong. Calibration on nominal examples alone cannot identify a posterior defect probability without information about defective observations and prevalence. The supervised committee used in Section 5 has a different training regime from the proposed nominal-only anomaly scorer.

2.4 | Synthesis and Gap

Table 1 contrasts selected approaches with the proposed AVID architecture. Its proposed row describes design capabilities rather than executed components. These comparisons are not an exhaustive claim that the integration is absent from prior work. The present evidence concerns a digit proxy with preview-assisted policies; learned stopping and asymmetric-cost operating points remain future evaluations.

Table 1. Capabilities of selected approaches and the proposed AVID architecture.

Method Family	Representative Works	View Selection	Learned Stopping	Epistemic Uncertainty	Asymmetric Costs	Budget-Aware Evaluation
Passive single-view detection	[4-6]	No	No	Partial	No	No
Classical active vision	[1], [2]	Yes, geometry-driven	No	No	No	No
Glimpse and attention policies	[11], [12]	Yes, learned	Implicit	No	No	Step cost only
Active learning	[16], [18]	Acquisition functions	Fixed budget	Yes	No	Partial
Sequential testing	[15]	No	Yes, optimal	No	Error rates	No
Proposed AVID	This work	Yes, evidence-driven	Yes	Yes	Yes	Yes

Note: The proposed row does not describe all components executed in the digit study.

3 | Problem Formulation and Research Questions

This section states the research questions and hypotheses, fixes the system model, and formalizes the decision problem. The formalization treats the inspection station as a measurement device whose configurations are actions and the part's defect state as the hidden variable. Notation is summarized in *Table 2*.

3.1 | Research Questions and Hypotheses

The design questions address evidence aggregation, planning over remaining budget, and exploration. They organize future tests of the complete architecture. The reported committee experiment evaluates handcrafted selection and stopping, rather than the learned policy or exploration procedure.

RQ1 asks how to aggregate available view evidence and quantify uncertainty under a finite acquisition budget.

RQ2 asks whether a learned joint policy improves the quality–cost trade-off at matched scorer capacity and information access.

RQ3 asks whether disagreement and information-based exploration improve learning. That learning comparison is unexecuted here.

H1 predicts a target detection AUC with fewer acquisitions for a learned policy under matched scorer and preview access.

H2 predicts an advantage for planning over myopic selection at small budgets. No reinforcement-learning planner is tested.

H3 predicts useful operating points under asymmetric decision costs. The reported symmetric payoff does not test it.

3.2 | System Model and Notation

A part x has a hidden label y∈{nominal,defective}. Acquiring configuration v∈V produces observation Zv, represented by an image Iv=T_v(x) with any specified sensing noise. M ensemble members produce scores and disagreement. The belief bt summarizes acquired evidence. The policy selects an unvisited view or a terminal decision d∈{accept,reject,escalate}, and acquires at most B views. Utility U(d,y) may encode

asymmetric decision costs in the design, while the executed experiment uses equal penalties for both error types.

Table 2. Notation used throughout the manuscript.

Symbol	Meaning
x, y	Part instance and true label in {nominal, defective}
V	Finite acquisition set of viewpoint-illumination configurations
I_v = T_v(x)	Image acquired under configuration v
h_1, ..., h_M	Ensemble of anomaly scorers
s_v, u_v	Ensemble mean score and disagreement for view v
b_t	Belief state aggregating acquisitions and evidence up to step t
d in {accept, reject, escalate}	Terminal decision actions
B, c_acq	Acquisition budget and per-acquisition cost
rho_FA, rho_FR, rho_E	False-accept, false-reject, and escalation penalties
U(d, y)	Terminal utility of Eq. (1)
EIG(v | b_t)	Expected information gain of acquiring v, Eq. (3)

3.3 | Formal Problem Statement

Eq. (1) defines the unshaped net-utility objective. *Eq. (2)* specifies an aggregation interface rather than a proven Bayesian update. *Eq. (3)* defines genuine expected information gain by averaging over the predictive distribution of a prospective observation Zv. It requires an observation model. Arbitrary entropy bonuses are not generally policy-invariant shaping; a decaying bonus alone does not guarantee convergence to an optimum of *Eq. (1)*. The experiment uses a preview entropy score instead of evaluating *Eq. (3)*.

$$\max_{\pi} \mathbb{E}\left[U(d_T, y) - \sum_{t=1}^{T} c_{acq}(v_t)\right], \quad T \le B. \tag{1}$$

$$b_{t+1} = \Psi(b_t, v_{t+1}, s_{t+1}, u_{t+1}), \quad s_v = \frac{1}{M}\sum_{m=1}^{M} h_m(I_v), \quad u_v = SD_m h_m(I_v). \tag{2}$$

$$EIG(v \mid b_t) = \mathbb{E}_{Z_v \sim p(\cdot \mid b_t, v)}\left[D_{KL}\left(p(y \mid b_t, Z_v, v) \| p(y \mid b_t)\right)\right]. \tag{3}$$

The finite budget and utility can price both errors and escalation. An ensemble score does not automatically constitute a calibrated belief; its interpretation must be validated. Formal expected information gain averages over observations before acquisition, whereas the tested rule evaluates actual candidate features in advance. Evaluation separates ordered-view AUC from deployed accuracy, acquisition count, and utility.

4 | Proposed Methodology

4.1 | Overview and Assumptions

The proposed design has an acquisition model, ensemble scorer, evidence aggregator, and learned policy over acquisition and terminal actions. Nominal-only anomaly scoring and PPO policy training are possible implementations, not the implemented committee experiment. Any ratio-based fallback requires a justified likelihood model for Wald error bounds; otherwise it is an interpretable heuristic.

The acquisition set is finite and known costs are specified per configuration.

Physical validity of transformed views requires measurements; two-dimensional resampling does not reproduce new viewpoints or relighting in general.

A nominal-only anomaly score requires additional calibration evidence before it is read as defect probability. The tested committee is supervised.

Prevalence and sample dependence must be specified. Drift may require recalibration but is not automatically corrected by it.

The acquisition budget B and terminal costs are externally specified.

4.2 | Acquisition Model and Evidence Extraction

The proposed acquisition model would apply geometric and photometric transformations, potentially using differentiable resampling [30]. Two-dimensional crops and resampling do not reproduce a new physical viewpoint or illumination geometry. The executed study uses only the six deterministic digit transforms of Section 5.1.

Each acquired view supplies member scores, their mean, and disagreement. The full architecture would use independently trained scorers [8]. Temperature scaling calibrates supervised probabilities under an appropriate validation design [10]; nominal-only calibration cannot determine defect odds by itself. Disagreement may indicate sensitivity among fitted members but need not increase for every informative or hazardous observation.

4.3 | Belief Encoding and Policy Architecture

The belief encoder consumes the set of evidence records together with their acquisition parameters. Each record is embedded and annotated with a positional code of its configuration, and an attention block with masking over unvisited configurations aggregates the records into a recurrent belief token [17]. The policy head outputs a categorical distribution over unvisited acquisitions and the three terminal actions, with visited views masked so that repeated acquisition is structurally impossible. A value head supports proximal policy optimization. The architecture is deliberately set-based: order invariance over acquisitions removes spurious sequence effects, and the mask makes budget accounting explicit, since the action space shrinks deterministically as the episode proceeds.

Two design choices deserve justification. First, acquisition parameters are embedded alongside evidence because the value of a view depends not only on what was seen but on what was looked at; two identical scores from orthogonal view geometries carry different implications. Second, the terminal actions are part of the same policy rather than a separate classifier, because the option value of continuing, one more acquisition versus deciding now, is exactly the quantity the value head must learn, and separating the decisions would force a hand-crafted arbitration between them.

4.4 | Training with Disagreement-Driven Exploration and Gain Shaping

A future PPO policy could optimize terminal utility minus acquisition costs [14]. Information and disagreement bonuses would define additional training objectives; their effects and decay schedule require empirical evaluation, without a policy-invariance guarantee. The ratio fallback can use target-error thresholds as design parameters, but formal sequential error bounds require verified likelihood assumptions and dependence treatment. Neither this training procedure nor escalation is executed by the committee study.

Algorithm 1. AVID decision loop for one part.

```
Input: part x; scorer ensemble h_1,…,h_M; evidence aggregator Ψ;
       policy π0; acquisition budget B; terminal decision costs;
       ratio thresholds A_L, A_U and a justified evidence model
E ← empty evidence set; visited ← empty set
for t = 1,…,B do
  b_t ← Ψ(E, visited)
  choose a_t from π0 given b_t and B−|visited|
  if a_t is a terminal decision d then
   compute a sequential evidence ratio under the specified model
   if the policy and ratio rule disagree, set d ← escalate
    emit d and the acquisition record; stop
  else
    I_v <- acquire(x, a_t);  (s_v, u_v) <- ensemble evidence
    E ← E ∪ {(a_t,s_v,u_v)}; visited ← visited ∪ {a_t}
  end if
end for
if no decision was emitted, apply the configured ratio rule;
  set d ← escalate if inconclusive; otherwise choose accept/reject, then emit d
Output: decision d in {accept, reject, escalate} and audit record
```

4.5 | Complexity, Scalability, and Failure Conditions

Per-view cost includes transformation, M scorer evaluations, evidence aggregation, and any sensor movement. Preview policies also score unvisited candidates. Such work is not covered by counting acquired views alone. Caching may reduce simulation work but does not remove a physical information-acquisition cost; runtime must be measured at the intended station.

Several failure conditions can be identified in advance. They are stated here so that negative results can be localized, and each is paired with its planned mitigation.

I. Transform-reality gap: two-dimensional transforms cannot capture parallax, true perspective, and surface-relief shading; conclusions must be re-validated on physical stations before transfer claims are made.

II. Ensemble degeneration: correlated members understate disagreement, starving both exploration and stopping; mitigated by member decorrelation and monitored by calibration metrics.

III. Cost misspecification: wrong penalties shift operating points silently; mitigated by sweeping the penalty ratio and reporting the full operating curve rather than a single point.

IV. Degenerate policies: always-reject or always-accept attractors under asymmetric costs; mitigated by curriculum design and by anchoring evaluation against the sequential-ratio baseline.

V. Prevalence drift: a shifted defect rate miscalibrates the posterior; mitigated by periodic recalibration and by reporting sensitivity to prevalence misspecification.

5 | Experimental Results and Analysis

5.1 | Setup: Corpus, Views, Committee, Policies, and Protocol

The reported study uses the 1,797 8×8 grayscale images distributed by scikit-learn's load_digits interface [31], derived from the UCI optical-recognition data [32] It defines proxy defects by digit classes ({3,8,9} in the main composition; prevalence 0.299), not industrial damage. Six transforms expose four quadrant crops and two global 2×2 block-mean downsamples, one preceded by Gaussian smoothing with σ=0.8. Training-split standardization is described per 16-feature block. Acquired blocks fill a 96-feature vector, and unacquired blocks are zero. MVTec evaluation [3] is future work, and a dated preregistration is not established.

The belief model is a supervised bagged committee of ten logistic regressions [33]. Bootstrap fitting includes random view masks, retaining each view with probability 0.5 and at least one view. The committee mean is

used as a defect score, without an automatic calibration guarantee. The five policies are random order, fixed order, maximum preview disagreement, a posterior-odds threshold rule labeled SPRT, and AVID preview selection with one-step stopping. AVID scores a candidate by current binary entropy minus the mean member entropy after inserting its actual features; this can be negative and is not the expected information gain of *Eq. (3)*. Its stopping score is six times the predicted reduction in min(p,1−p), compared with cost 0.15. The ratio baseline uses posterior odds divided by prior odds between thresholds 0.05 and 20. Overlapping deterministic views do not justify treating this discriminative ratio as a product of independent likelihood ratios, so Wald guarantees are not assigned. Confidence baselines use threshold 0.90. These settings are reported without a timestamped prespecification record.

AVID and disagreement-greedy selection evaluate actual candidate features before their acquisition is charged. This is an oracle preview in the proxy experiment, not a measured coarse physical preview. Non-preview baselines have unequal information access. Net utility is correct=+1, wrong=−5, minus 0.15 per acquired view; its expectation is 1−6e−0.15v, where e is the error rate and v the mean count. Changing a wrong decision into a correct one gains six, not four. The protocol uses stratified 60/20/20 splits, twenty seeds for each of three compositions, and separate no-stopping AUC curves and deployed stopping metrics. The stated 126 s runtime and software configuration require original logs; view counts omit preview scoring and possible physical preview costs.

5.2 | Main Comparison (RQ1)

Table 3 reports one-view AUC 0.9522 for preview-based AVID, versus 0.8636 for random, 0.8556 for fixed, and 0.7541 for preview-disagreement selection. At two views the AVID value is 0.9667, and all policies reach 0.9806 at six. This comparison mixes selection with access to unacquired features. The lower disagreement-greedy result is an observation, not proof that disagreement always selects the least informative view.

AVID stops after 1.65 views on average versus 3.90 for the ratio baseline, a reduction of about 58%. Its reported accuracy is 0.8650 and utility −0.058, versus 0.9368 and +0.036 for the ratio baseline. Fixed and random confidence stopping have utilities +0.198 and +0.190. The source primary utility contrast reports W=11, p≈1.0e−4. Thus lower counted acquisition volume does not imply better utility under the stated payoff. Inference remains exploratory across overlapping repeated splits.

Table 3. Main-composition results over twenty repeated splits (mean±SD where applicable).

Acquisition Policy	AUC @1 View	AUC @2 Views	AUC @6 Views	Mean Views	Accuracy	Cost-Adjusted Utility
Random order	0.8636	0.9261	0.9806	2.92 ± 0.08	0.9382 ± 0.0097	0.190 ± 0.062
Fixed sequence	0.8556	0.9279	0.9806	2.67 ± 0.11	0.9331 ± 0.0127	0.198 ± 0.076
Uncertainty sampling	0.7541	0.8554	0.9806	4.57 ± 0.08	0.9346 ± 0.0093	-0.078 ± 0.057
Ratio-threshold heuristic (source label: SPRT)	0.8568	0.9221	0.9806	3.90 ± 0.09	0.9368 ± 0.0098	0.036 ± 0.063
AVID (proposed)	0.9522	0.9667	0.9806	1.65 ± 0.05	0.8650 ± 0.0146	-0.058 ± 0.085

Note: Ordered-view AUC disables early stopping; counts, accuracy, and utility enable the stopping rules. Preview access differs across policies.

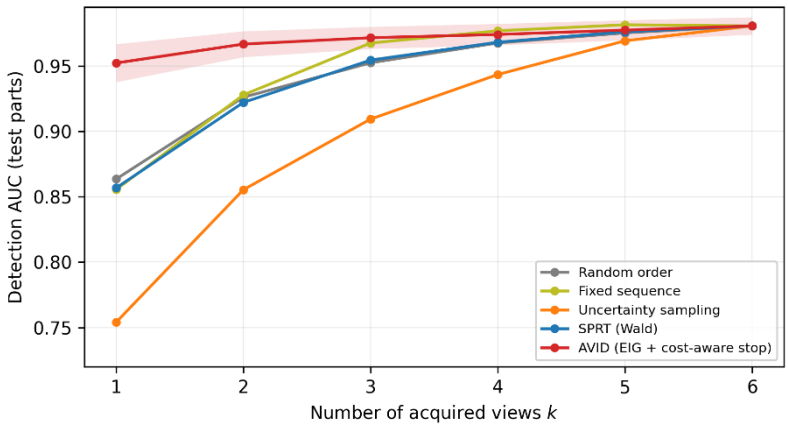


Fig. 1. Detection AUC versus counted acquired views (main composition, mean over twenty splits; AVID band ±SD).

The AVID and disagreement policies use actual unacquired candidate features as previews (*Fig. 1*).

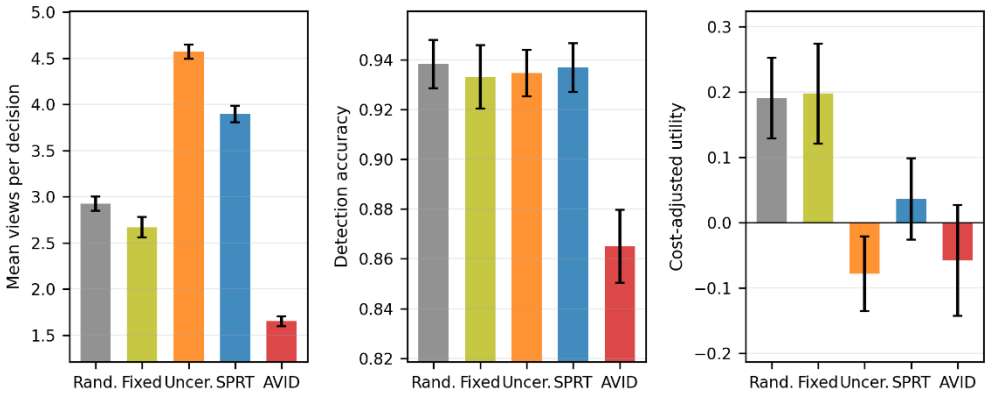


Fig. 2. Deployed behaviour of the five policies over twenty seeds (mean ± one standard deviation): Mean views per decision, detection accuracy, and cost-adjusted utility U = correct - 5 x wrong - 0.15 x views.

5.3 | Ablation (RQ2)

Table 4 shows the accuracy–cost trade-off. Always acquiring six views gives accuracy 0.9392 and utility −0.265, compared with 0.8650 and −0.058 for AVID stopping. The utility gain ≈0.207 combines cost saving $0.15\times(6-1.65)=0.6525$ with loss in decision payoff $6\times(0.9392-0.8650)=0.4452$. Accuracy is therefore materially lower, not almost unchanged. Replacing preview selection with random next views gives 3.48 views, accuracy 0.9065, and utility −0.083. These contrasts concern combined selection/stopping behavior and preview access.

Table 4. AVID ablations on the main defect composition (twenty seeds).

AVID Variant	Cost-Adjusted Utility	Mean Views	Accuracy
Full AVID (preview entropy selection + cost-aware stopping)	-0.058 ± 0.085	1.65	0.8650
− Preview selection (random next view, same stopping)	-0.083 ± 0.089	3.48	0.9065
− Cost-aware stopping (always 6 views)	-0.265 ± 0.055	6.00	0.9392

5.4 | Robustness: Defect-Class Composition (RQ3)

Under the thin-stroke composition, AVID reports 1.32 views and utility −0.243; under the alternative dense composition its utility is approximately zero. Composition changes task labels, prevalence, and view informativeness together. These patterns do not isolate a causal effect of semantic evidence sparsity or validate reinforcement-learning exploration.

Table 5. Robustness across defect-class compositions (twenty seeds each; best baseline among random order, fixed sequence, and uncertainty sampling).

Defect Classes	Prev.	AVID Views	AVID Utility	SPRT Views	SPRT Utility	Best Baseline (Utility)
{3, 8, 9} (dense strokes)	0.299	1.65	-0.058 ± 0.085	3.90	0.036 ± 0.063	Fixed sequence (0.198 ± 0.076)
{1, 7} (thin strokes)	0.201	1.32	-0.243 ± 0.047	4.02	-0.032 ± 0.069	Random order (0.168 ± 0.085)
{0, 6, 8} (alternative dense)	0.297	1.66	-0.000 ± 0.078	4.01	0.005 ± 0.059	Fixed sequence (0.160 ± 0.051)

Note: The source calls the main composition pre-registered, but supplies no dated protocol. The SPRT columns denote the tested posterior-ratio threshold heuristic.

5.5 | Statistical Analysis

Table 6 retains the source Wilcoxon [34], Holm [35], and Friedman [36], [37] summaries. The reported Friedman statistic is $\chi^2(4)=69.72$, $p=2.60e{-}14$. Utility mean ordering is fixed, random, ratio, AVID, then disagreement. A nonsignificant AVID-versus-disagreement test is not equivalence. The twenty repeated splits reuse the same corpus with overlapping examples, so the test assumptions and dependence require inspection. Wilcoxon assesses signed-rank differences, not arithmetic mean direction or acquisition-cost efficiency.

Table 6. Statistical analysis (main composition, twenty seeds).

Contrast (Main Composition)	Statistic W	p	p (Holm)
Primary: AVID vs SPRT, cost-adjusted utility	11.0	1.05e-04	—
AVID vs Random order, cost-adjusted utility	0.0	1.91e-06	7.63e-06
AVID vs Fixed sequence, cost-adjusted utility	0.0	1.91e-06	7.63e-06
AVID vs Uncertainty sampling, cost-adjusted utility	83.0	4.30e-01	4.30e-01
AVID vs SPRT, mean views per decision	0.0	1.91e-06	—
Friedman across the five policies (utility)	$\chi^2(4) = 69.72$	2.60e-14	—

5.6 | Discussion: Findings, Honest Losses, and Scope

The preview selector has the highest displayed one-view AUC, but it uses privileged candidate content. The reported 78% of 48.6 wrong decisions per seed that would become correct with all six views is a retrospective overlap count requiring prediction records. It identifies potentially recoverable exits, not their unique cause. Complementary-view value is a possible explanation for myopic stopping, but confidence errors, preview modeling, and split variability are alternatives. Full-view accuracy 0.9392 does not imply that every individual additional view is valuable.

The committee realization has lower AVID utility than the ratio and confidence-stopping baselines, despite fewer counted views. The proxy uses deterministic 8×8 images and supervised linear models. Reinforcement learning, attention aggregation, escalation, and asymmetric error costs are not executed, so H1–H3 about the complete design are not resolved by this campaign.

The reported proof-of-concept outcomes require the exact split indices, view transforms, committee predictions, policy traces, and executable implementation for independent verification.

6 | Discussion: Limitations and Counterarguments

6.1 | Limitations

The digit proxy does not measure industrial damage, real viewpoint costs, or sensing noise. Its results cannot establish either absolute or relative performance at a physical station. Nominal-only scoring is a proposed component, while the tested committee uses defect labels. Unpriced oracle preview is a central information-access limitation. Repeated splits, probability calibration, joint-view dependence, and prevalence changes further bound inference. Station scheduling and asymmetric economics remain untested.

6.2 | Alternative Viewpoints and Responses

Acquiring all views may be rational when error costs dominate acquisition costs. Here it improves accuracy while lowering utility under one chosen payoff. Greedy and confidence rules are useful comparators, but this experiment contains no learned planner. A ratio rule is auditable without being a formally valid Wald test. Physical previews and escalation need explicit observation and cost models.

7 | Conclusion and Future Directions

The reported digit comparison separates view ordering from stopping outcomes. Preview-based AVID reaches single-view AUC 0.9522, but receives unacquired candidate content before cost accounting. Its stopping rule uses 1.65 views and accuracy 0.8650, versus 3.90 and 0.9368 for the ratio baseline; net utility is lower, −0.058 versus +0.036. Always six views yields accuracy 0.9392 at a higher acquisition cost. These results show a scoped accuracy–cost trade-off under oracle preview. They do not support a general recommendation to use the selector for physical acquisition or an interpretation of the stopping rule as an optimality bound.

Further evaluation should share or price preview access, train and test any learned multi-step policy, calibrate probabilities using independent data, and report sensitivity to costs and prevalence. Industrial multi-view data and a justified sequential observation model are required before transfer, asymmetric-cost, or Wald-guarantee claims.

Author Contributions

Sh. K.: conceptualization, methodology, software, and writing—original draft. M. P.: data curation, formal analysis, and visualization. M. P.: validation, supervision, and writing—review and editing. All authors read and approved the final manuscript.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Data Availability

The digit corpus is publicly distributed through scikit-learn and originates from the UCI optical-recognition data. The study-specific split indices and raw outputs are required for reproducibility; an archived identifier is not specified.

Code Availability

A public repository and archived version for the study-specific implementation are not specified.

Conflicts of Interest

The authors declare no conflicts of interest.

Generative AI disclosure

Drafting assistance was provided by a generative AI system under the authors' direction. The authors remain responsible for verifying scientific content, numerical provenance, and the accuracy of the disclosure.

References

[1] Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE*, *76*(8), 966–1005. https://doi.org/10.1109/5.5968

[2] Aloimonos, J., Weiss, I., & Bandyopadhyay, A. (1988). Active vision. *International journal of computer vision*, *1*(4), 333–356. https://doi.org/10.1007/BF00133571

[3] Bergmann, P., Fauser, M., Sattlegger, D., & Steger, C. (2019). MVTec ad—a comprehensive real-world dataset for unsupervised anomaly detection. *2019 IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 9584–9592). IEEE. https://doi.org/10.1109/CVPR.2019.00982

[4] Akcay, S., Atapour-Abarghouei, A., & Breckon, T. P. (2018). Ganomaly: semi-supervised anomaly detection via adversarial training. *Asian conference on computer vision* (pp. 622–637). Springer. https://arxiv.org/abs/1805.06725

[5] Defard, T., Setkov, A., Loesch, A., & Audigier, R. (2021). Padim: a patch distribution modeling framework for anomaly detection and localization. *International conference on pattern recognition* (pp. 475–489). Springer. https://arxiv.org/abs/2011.08785

[6] Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., & Gehler, P. (2022). Towards total recall in industrial anomaly detection. *2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 14298–14308). IEEE. https://doi.org/10.1109/cvpr52688.2022.01392

[7] Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., … & Müller, K.-R. (2021). A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, *109*(5), 756–795. https://doi.org/10.1109/jproc.2021.3052449

[8] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, *30*. https://doi.org/10.48550/arXiv.1612.01474

[9] Gal, Y., & Ghahramani, Z. (2016). Dropout as a bayesian approximation: representing model uncertainty in deep learning. *International conference on machine learning* (pp. 1050–1059). PMLR. https://doi.org/10.48550/arXiv.1506.02142

[10] Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. *International conference on machine learning* (pp. 1321–1330). PMLR. https://proceedings.mlr.press/v70/guo17a.html

[11] Mnih, V., Heess, N., Graves, A., & Kavukcuoglu, K. (2014). *Recurrent models of visual attention* (Vol. 27). https://arxiv.org/abs/1406.6247

[12] Ba, J., Mnih, V., & Kavukcuoglu, K. (2014). *Multiple object recognition with visual attention*. https://doi.org/10.48550/arXiv.1412.7755

[13] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., & Bellemare, M. G. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. https://doi.org/10.1038/nature14236

[14] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). *Proximal policy optimization algorithms*. https://arxiv.org/abs/1707.06347

[15] Wald, A. (1992). Sequential tests of statistical hypotheses. In *Breakthroughs in statistics: Foundations and basic theory* (pp. 256–298). Springer. https://doi.org/10.1214/aoms/1177731118

[16] Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The annals of mathematical statistics*, *27*(4), 986–1005. https://doi.org/10.1214/aoms/1177728069

[17] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., … & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems* (Vol. 2017-Decem, pp. 5999–6009). NIPS. https://arxiv.org/abs/1706.03762

[18] Settles, B. (2009). *Active learning literature survey*. https://burrsettles.com/pub/settles.activelearning.pdf

[19] Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., & Steger, C. (2021). The MVTec anomaly detection dataset: a comprehensive real-world dataset for unsupervised anomaly detection. *International journal of computer vision*, *129*(4), 1038–1059. https://doi.org/10.1007/s11263-020-01400-4

[20] Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., & Steger, C. (2022). Beyond dents and scratches: Logical constraints in unsupervised anomaly detection and localization. *International journal of computer vision*, *130*(4), 947–969. https://doi.org/10.1007/s11263-022-01578-9%0A%0A

[21] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778). IEEE. https://doi.org/10.1109/cvpr.2016.90

[22] Simonyan, K., & Zisserman, A. (2014). *Very deep convolutional networks for large-scale image recognition*. https://arxiv.org/abs/1409.1556

[23] Hendrycks, D., & Dietterich, T. (2019). *Benchmarking neural network robustness to common corruptions and perturbations*. https://arxiv.org/abs/1903.12261

[24] Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, *101*(1–2), 99–134. https://doi.org/10.1016/s0004-3702(98)00023-x

[25] Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: off-policy maximum entropy deep reinforcement learning with a stochastic actor. *International conference on machine learning* (pp. 1861–1870). PMLR. https://proceedings.mlr.press/v80/haarnoja18b

[26] Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine learning*, *47*(2), 235–256. https://doi.org/10.1023/a:1013689704352

[27] Pathak, D., Agrawal, P., Efros, A. A., & Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. *International conference on machine learning* (pp. 2778–2787). PMLR. https://proceedings.mlr.press/v70/pathak17a.html?ref=https://githubhelp.com

[28] Burda, Y., Edwards, H., Storkey, A., & Klimov, O. (2018). *Exploration by random network distillation*. https://doi.org/10.48550/arXiv.1810.12894

[29] Elkan, C. (2001). *The foundations of cost-sensitive learning* [presentation]. International joint conference on artificial intelligence (Vol. 17, pp. 973–978). https://cseweb.ucsd.edu/~elkan/rescale.pdf

[30] Jaderberg, M., Simonyan, K., & Zisserman, A. (2015). *Spatial transformer networks*. https://arxiv.org/abs/1506.02025

[31] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., & Grisel, O. (2011). Scikit-learn: Machine learning in Python. *The journal of machine learning research*, *12*, 2825–2830. https://jmlr.org/papers/v12/pedregosa11a.html

[32] Alpaydin, E., & Kaynak, C. (1998). *Optical recognition of handwritten digits*. https://doi.org/10.24432/C50P49

[33] Breiman, L. (1996). Bagging predictors. *Machine learning*, *24*(2), 123–140. https://doi.org/10.1023/a:1018054314350

[34] Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics bulletin*, *1*(6), 80–83. https://doi.org/10.2307/3001968

[35] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, *6*(2), 65–70. https://www.jstor.org/stable/4615733

[36] Friedman, M. (1937). The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American statistical association*, *32*, 675–701. https://doi.org/10.1080/01621459.1937.10503522

[37] Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of machine learning research*, *7*(1), 1–30. http://www.jmlr.org/papers/volume7/demsar06a/demsar06a.pdf