Coverage-Optimal Robotic Colonoscopy via Patient-Specific Digital-Twin World Models

Luca Carlini*, Alberto Rota*, Sofia Breschi, Laura Cruciani, Matteo Di Mauro, Alessandra Marasi,
Nicola Ancellotti, Leonardo Passoni, Elena De Momi

Department of Electronics, Information and Bioengineering , Politecnico di Milano

*These authors contributed equally.

Abstract—Post-colonoscopy interval cancers account for 2–8% of colorectal cancers and are driven primarily by mucosa that goes unexamined during withdrawal. The Adenoma Detection Rate (ADR), the strongest quality indicator, is inversely and dose-dependently associated with interval-cancer risk, yet current robotic and autonomous colonoscopy systems are optimized for reliably reaching the caecum rather than for maximizing patient-specific mucosal coverage. We propose a framework that segments a patient CT colonography scan into a patient-specific *digital twin* of the colon, which serves as the environment in which a navigation agent learns. Within the twin, a *world model* predicts the next state (the endoluminal RGB view, the camera pose, and the cumulative coverage) resulting from each candidate motion, and a navigation policy is trained in a closed learning loop against a reward that credits novel coverage gain weighted by view quality (viewing distance and incidence angle). The trained policy then drives a robotic shared-autonomy deployment on the live endoscopic feed and is evaluated by mucosal coverage, ADR, and number of missed lesions. By producing per-patient blind-spot maps and coverage-optimal inspection plans, the approach standardizes inspection quality, reduces operator dependence, and directly targets the root cause of interval cancers. In contrast to systems optimized for reaching the caecum, it reframes autonomous colonoscopy around the quantity that interval-cancer risk actually depends on: how completely and how well the mucosa is seen.

Keywords—surgical robotics, colonoscopy, reinforcement learning, world models, digital twin, autonomous navigation.

1. Clinical and Technological Need

Colonoscopy is the reference test for colorectal-cancer screening, yet 2–8% of colorectal cancers are post-colonoscopy interval malignancies that emerge after an examination in which the responsible lesion was present but not detected [1]. The Adenoma Detection Rate (ADR) is the strongest validated quality indicator and is inversely and dose-dependently linked to interval-cancer risk: risk decreases by roughly 3% per ADR point and is more than tenfold higher below a 20% ADR [2, 3]. A primary and modifiable cause of these misses is mucosa that is never adequately visualized during withdrawal, typically hidden behind haustral folds or at the flexures. Recent robotic and autonomous colonoscopy systems have made substantial progress on locomotion and on safely reaching the caecum [3, 4], but they optimize for reliable insertion rather than for how much of the colonic surface is actually inspected. Inspection completeness therefore remains largely operator-dependent and unmeasured, motivating a system that explicitly optimizes patient-specific mucosal coverage.

2. Methodology

The proposed pipeline is summarized in Fig. 1. A patient's CT colonography scan is segmented into a patient-specific *digital twin*: a three-dimensional model of that individual's colon that serves as the environment in which the navigation agent learns. Because the twin provides the full surface geometry, the fraction of mucosa observed along a given trajectory can be computed exactly, turning inspection completeness into a measurable, optimizable quantity. Inside the twin, a *world model* predicts the future state

$$\mathbf{S}(t) = \left[\text{RGB},\ \mathbf{T},\ C_{\%}\right]_t,$$

comprising the simulated endoluminal RGB view, the camera pose $\mathbf{T}$, and the cumulative coverage $C_{\%}$, that would result from executing each candidate action $\mathbf{A}(t) = [\mathbf{R} \mid \mathbf{t}]_t$. Predicting future states lets the agent anticipate the consequences of its motions and plan toward unseen regions rather than react frame by frame. A navigation policy proposes actions and is trained in a closed learning loop against the policy reward

$$R(\theta) = \Delta C_{\%} \cdot q(d, \alpha),$$

which credits only the *novel* coverage gain $\Delta C_{\%}$ obtained at each step and weights it by a view-quality term q that decreases for unfavorable viewing geometry, i.e. a large viewing distance d or an oblique incidence angle α. Mucosa seen only briefly, from too far, or at a grazing angle therefore contributes little, encouraging trajectories that both reach new surface and inspect it well. The trained policy is then transferred to a robotic *shared-autonomy* deployment, where it guides camera motion on the live endoscopic feed with the clinician remaining in the loop.

3. Evaluation and Expected Impact

The framework is assessed with metrics that map directly onto the clinical failure mode it targets: mucosal coverage percentage quantifies inspection completeness, while ADR and the number of missed lesions quantify detection performance. Beyond a single score, the digital twin yields per-patient blind-spot maps

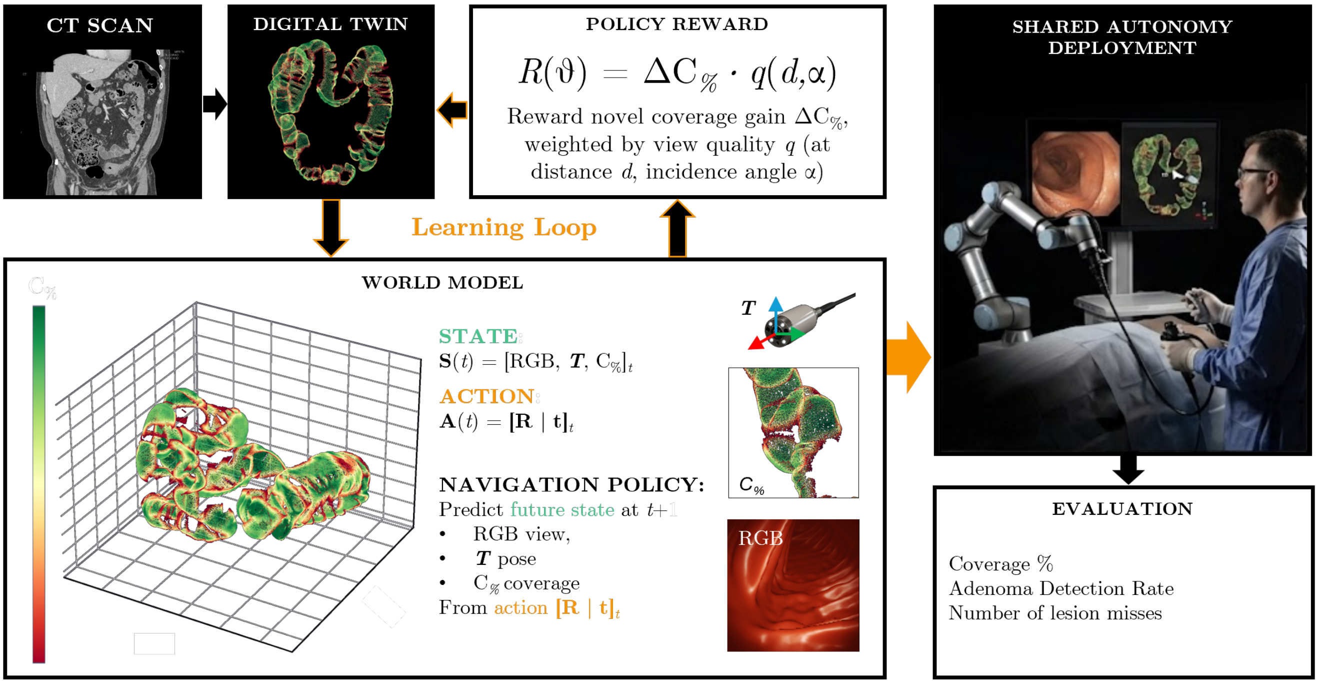

Figure 1: Overview of the proposed pipeline. A patient CT colonography scan is segmented into a patient-specific digital twin of the colon. Inside the twin a world model predicts the next state $\mathbf{S}(t) = [\mathrm{RGB}, \mathbf{T}, C_{\%}]_t$ from each action $\mathbf{A}(t) = [\mathbf{R} \mid \mathbf{t}]_t$, and a navigation policy is trained in a closed learning loop against the coverage reward $R(\theta) = \Delta C_{\%} \cdot q(d, \alpha)$. The learned policy then drives a robotic shared-autonomy deployment on the live endoscopic feed, evaluated by coverage, ADR and number of missed lesions.

that localize the regions hardest to visualize, together with coverage-optimal inspection plans tailored to each colon's anatomy. By making coverage explicit and optimizing it directly, the approach standardizes inspection quality across operators, reduces the dependence of outcomes on individual skill, and attacks the root cause of interval cancers: mucosa that is never examined.

4. Future Work

Two directions are planned. First, integration with physical robotic colonoscopy platforms, moving from in-twin training to semi-autonomous and ultimately fully autonomous coverage optimization during real procedures. Second, uncertainty-aware mapping: by quantifying the confidence of the coverage estimate across the surface, the system can flag low-confidence regions and prioritize their re-inspection, further reducing the chance that mucosa is silently left unseen.

References

[1] M. Finocchiaro *et al.*, "Clinically validated dataset of 435 human colons segmented from CT colonography," *Sci. Data*, 2026.

[2] M. F. Kaminski *et al.*, "Quality indicators for colonoscopy and the risk of interval cancer," *N. Engl. J. Med.*, 2010.

[3] J. W. Martin *et al.*, "Enabling the future of colonoscopy with intelligent and autonomous magnetic manipulation," *Nat. Mach. Intell.*, 2020.

[4] D. Orsi *et al.*, "Constrained reinforcement learning and formal verification for safe colonoscopy navigation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2023.

Coverage-Optimal Robotic Colonoscopy via Patient-Specific Digital-Twin World Models

Luca Carlini*, Alberto Rota*, Sofia Breschi, Laura Cruciani, Matteo Di Mauro, Alessandra Marasi, Elena De Momi

POLITECNICO MILANO 1863

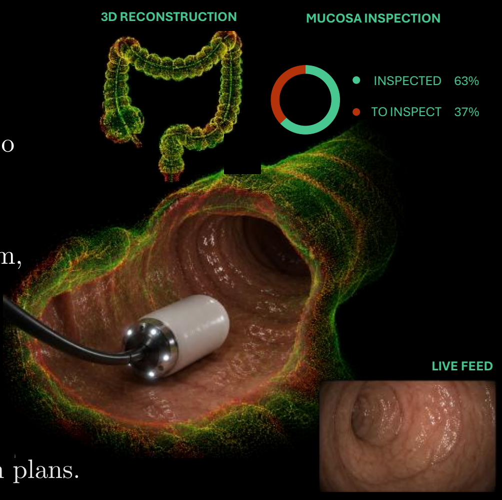

Clinical and Technology Need

- **2** to **8%** of colorectal cancers are post-colonoscopy interval cancers. [1].
- Adenoma Detection Rate (ADR) is inversely, dose-dependently linked to interval cancer, roughly **3% fewer** per ADR point, with over ten times the risk below a 20% ADR. [2,3]
- Robotic and autonomous colonoscopy, optimises for reaching the caecum, not for patient-specific coverage.

Impact

- Standardises quality and **reduces operator dependence.**
- Yields **per-patient blind-spot maps** and coverage-optimal inspection plans.
- Targets the root cause of interval cancers by systematically **minimising unexamined mucosa.**

Methodology

A patient **CT SCAN** is segmented into a patient-specific **DIGITAL TWIN** of the colon, which serves as the environment the agent learns in. Inside the twin, a **WORLD MODEL** predicts the next **STATE** [RGB, pose $\boldsymbol{T}$, coverage $C_{\%}$] from each **ACTION** [R | t]; the NAVIGATION POLICY proposes actions and the Learning Loop trains it against the **POLICY REWARD** R, rewarding coverage gain and view fidelity. The trained policy then drives **SHARED AUTONOMY DEPLOYMENT** on the live feed, with **EVALUATION** by coverage %, adenoma detection rate and lesion misses.

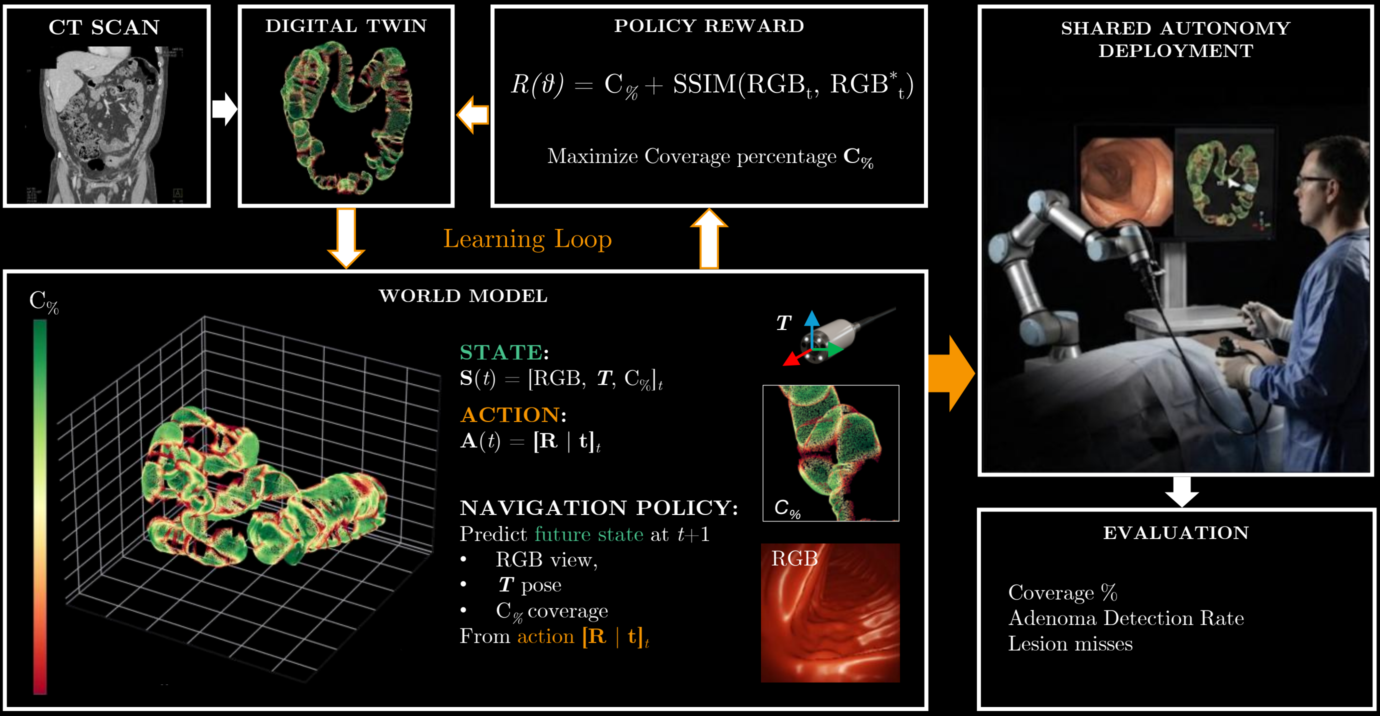

CT SCAN

DIGITAL TWIN

POLICY REWARD

$$R(\vartheta) = C_{\%} + \mathrm{SSIM}(\mathrm{RGB}_t, \mathrm{RGB}^*_t)$$

Maximize Coverage percentage $\mathbf{C}_{\%}$

Learning Loop

WORLD MODEL

STATE:
$\mathbf{S}(t) = [\mathrm{RGB}, \boldsymbol{T}, C_{\%}]_t$

ACTION:
$\mathbf{A}(t) = [\mathbf{R} \mid \mathbf{t}]_t$

NAVIGATION POLICY:
Predict future state at $t+1$

- RGB view,
- $\boldsymbol{T}$ pose
- $C_{\%}$ coverage

From action $[\mathbf{R} \mid \mathbf{t}]_t$

SHARED AUTONOMY DEPLOYMENT

EVALUATION

Coverage %
Adenoma Detection Rate
Lesion misses

Future Work

- Integration with robotic colonoscopy platforms for semi-autonomous or autonomous coverage optimization.
- Uncertainty-aware mapping to identify regions with low confidence and prioritize re-inspection of potentially missed mucosa.

[1] M. Finocchiaro *et al.*, "Clinically validated dataset of 435 human colons segmented from CT colonography," *Sci. Data*,, 2026.
[2] M. F. Kaminski *et al.*, "Quality indicators for colonoscopy and the risk of interval cancer," *N. Engl. J. Med.*,, 2010.
[3] J. W. Martin *et al.*, "Enabling the future of colonoscopy with intelligent and autonomous magnetic manipulation," *Nat. Mach. Intell.*,, 2020.
[4] D. Orsi *et al.*, "Constrained reinforcement learning and formal verification for safe colonoscopy navigation," in *Proc. IEEE/RSJ IROS*, 2023.