

1 Transferable Deep Reinforcement Learning with Edge-Contour-Depth 2 Fusion for Autonomous Wireless Capsule Endoscopy Navigation

3

4 *Haoxuan Wu^{1†}, Haitao Gao^{2†}, Qingyang Liu^{1†}, Sishen Yuan^{1*}, Haiyang Fang³, Mingwu Su¹,*
5 *Baijia Liang¹, Yongzun Yang¹, Long Bai¹, Wenzhen Dong¹, Dihong Xie¹, Shijian Su¹, Jiewen*
6 *Lai¹, Shing Shin Cheng³, Zhen Li⁴, Xiuli Zuo⁴, Hongliang Ren^{1*}*

7

8 H. Wu, Q. Liu, S. Yuan, M. Su, B. Liang, Y. Yang, L. Bai, W. Dong, D. Xie, S. Su, J. Lai, H. Ren

9 Department of Electronic Engineering, The Chinese University of Hong Kong, Hong Kong SAR, China

10 hxwu@link.cuhk.edu.hk; sishenyuan@cuhk.edu.hk; hren@cuhk.edu.hk

11

12 H. Gao

13 School of Computer Science and Engineering, University of New South Wales, Sydney, NSW 2052, Australia

14

15 H. Fang, S. S. Cheng

16 Department of Mechanical and Automation Engineering, The Chinese University of Hong Kong, Hong Kong
17 SAR, China

18

19 Z. Li, X. Zuo

20 Qilu Hospital of Shandong University, Shandong Province, China

21

22 † These authors contributed equally to this work.

23

24 Keywords: wireless capsule endoscopy(wce), deep reinforcement learning, model transferability, sim-to-real
25 transfer, autonomous wce navigation

26

27 **Abstract**

28 Wireless capsule endoscopy (WCE) enables painless, minimally invasive visualization of the
29 gastrointestinal tract. Still, its diagnostic potential is limited by incomplete mucosal coverage
30 and poor transferability of existing navigation methods across patient anatomies. We propose
31 a transferable, anatomical landmark-guided deep reinforcement learning framework for robust
32 autonomous gastric navigation. Leveraging a lightweight edge-contour-depth fusion module,
33 our policy operates on stable, low-dimensional landmark coordinates rather than high-
34 dimensional video streams. This design effectively bridges the sim-to-real visual gap and
35 ensures robustness across diverse anatomies, enabling low-cost deployment by reducing

36 computational overhead. In simulations across eight patient-derived models, the method
37 achieves >97% coverage within 50 s, significantly outperforming vanilla Proximal Policy
38 Optimization, Soft Actor-Critic, and Deep Q-Network agents by enhancing coverage and
39 minimizing variance. To ensure deployment reliability, a two-stage sim-to-real pipeline
40 supported by an adaptive dynamic programming controller actively mitigates physical
41 disturbances, including actuator latency and peristalsis. Ex-vivo experiments across five
42 independent scans demonstrate high coverage stability, achieving a mean coverage of 87% and
43 a 53% reduction in procedure time compared with expert manual control. This study establishes
44 a scalable paradigm for autonomous, high-coverage endoscopic navigation, advancing the
45 clinical deployment of intelligent WCE systems for GI diagnostics.

46

47 **Introduction**

48 Wireless capsule endoscopy (WCE) has revolutionized gastrointestinal diagnostics by enabling
49 comprehensive, minimally invasive visualization of the entire digestive tract. Unlike
50 conventional approaches, WCE provides unprecedented access to the small intestine and
51 stomach, facilitating the early detection of mucosal lesions and subtle pathological changes
52 that might otherwise be missed¹⁻⁴. Among these, gastric polyps—sessile or pedunculated
53 protrusions arising from the gastric epithelium or submucosa—are of particular clinical
54 concern due to their malignant potential⁵⁻⁷. The risk of malignant progression is closely tied to
55 histological subtype, underscoring the importance of thorough, systematic mucosal inspection
56 during gastric examination^{8,9}. Maximizing mucosal surface coverage is therefore a central
57 requirement for next-generation WCE systems, with direct implications for diagnostic
58 reliability and patient outcomes¹⁰.

59 Achieving complete gastric coverage, however, remains challenging due to the absence of fully
60 autonomous navigation. Current practice relies on manual capsule manipulation within a highly
61 complex and dynamically deforming gastric lumen. Navigation is typically achieved through
62 externally applied magnetic fields, generated via permanent magnets or electromagnetic coils
63 placed around the patient¹¹⁻¹³. Early methods, such as hand-held magnets, offered low precision
64 and were heavily dependent on the operator's skill and judgement¹⁴. Subsequent technological
65 advances have included robotic-arm assistance¹⁵⁻¹⁷, joystick-controlled electromagnetic
66 platforms offering greater degrees of freedom (DOF)¹⁸, and high-DOF serial-link manipulators
67 for real-time control in fluid-distended stomachs¹⁹⁻²¹. While these approaches have improved
68 positioning accuracy and coverage fidelity²², they still suffer from incomplete automation,
69 restricted DOF, and the inherent complexity of navigating within an unstructured, deformable
70 organ²³. Consequently, the lack of robust autonomous capabilities continues to limit consistent,
71 operator-independent gastric coverage.

72 A further bottleneck is real-time localization and coverage assessment. Physicians must infer
73 capsule position and coverage status from limited endoscopic imagery, often collected under
74 low illumination and with ambiguous anatomical context²⁴⁻²⁶. Recent deep learning-based
75 localization methods that integrate visual, temporal, and motion cues have improved
76 anatomical landmark identification²⁷⁻²⁹, but their performance depends heavily on specific

77 organ geometries, device optics, and dataset diversity, limiting transferability across patient
78 anatomies.

79 Deep reinforcement learning (DRL) has shown promise for autonomous capsule trajectory
80 tracking and target reaching^{30,31}. However, existing studies adopt a conventional “end-to-end”
81 visual DRL paradigm, training policies by feeding high-dimensional endoscopic images
82 directly into the agent’s observation space. Although this approach can achieve high coverage
83 when evaluated within its original training environment^{32,33}, this design introduces several
84 critical limitations. First, it causes policies to overfit to the specific morphological and textural
85 features of the training environment, which severely restricts scalability and reduces
86 transferability to unseen geometries. Second, this reliance on raw visual input introduces two
87 further barriers to clinical translation: (1) computational demands for training and deployment,
88 which are infeasible for low-cost clinical hardware, and (2) vulnerability to the reality gap and
89 visual variations across patients, leading to poor transferability. Since comprehensive patient-
90 specific training datasets are impractical to obtain, improving transferability by overcoming
91 the limitations of visual-based DRL is essential for clinical deployment.

92 Here, we present a unified framework that integrates principled navigation feature selection,
93 state-of-the-art DRL, and robust optimal control to achieve reliable and transferable gastric
94 coverage. Based on the key insight that while mucosal imagery is highly variable, gastric
95 anatomical structures are highly conserved across patients³⁴. We therefore establish explicit
96 criteria—universality, distinctiveness, and navigational utility—for selecting anatomical
97 landmarks. This principled approach allows us to replace volatile, high-dimensional visual data
98 with stable, low-dimensional landmark coordinates, enabling DRL agents to consistently
99 interpret spatial cues across anatomically diverse stomachs. Building on this foundation, we
100 systematically train and benchmark PPO, SAC, and DQN for autonomous discovery of
101 coverage-maximizing trajectories in complex gastric environments. Our approach circumvents
102 the challenges of transferability and high computational cost by training a lightweight DRL
103 agent whose policy relies not on raw visual data, but on stable, low-dimensional spatial
104 coordinates—representing salient visual features and depth cues—extracted via our
105 lightweight edge-contour-depth fusion module. This design enhances policy transferability and
106 preserves high performance in anatomically varied and unseen stomach models. To close the
107 sim-to-real gap, we further design an optimal control scheme that compensates for actuator
108 latency, sensor noise, and other real-world disturbances, enabling direct policy deployment on
109 physical WCE platforms with accurate, rapid trajectory execution^{35,36}. Thus, these
110 contributions advance autonomous endoscopic navigation toward clinically viable systems
111 capable of consistent, high-coverage operation across diverse anatomical and procedural
112 scenarios (Figure 1c).

113

114 **Results**

115 **Design of Offline AL-DRL Framework**

116 The offline anatomical landmark-guided deep reinforcement learning (AL-DRL) framework
117 adopts a monocular camera-equipped WCE as the indirectly actuated robotic platform. It
118 integrates a supporting hardware system, a dual-modal perception module combining visual
119 and six-degree-of-freedom pose information, and a decision-making module (Figure 1). The
120 hardware system consists of two primary components: an internal section, containing the
121 single-camera WCE, which captures gastrointestinal images and is equipped with an array of
122 internal permanent magnet (IPM) for magnetic control; an external section, comprising a
123 6-DOF robotic arm (UR5, Universal Robots) with an external permanent magnet (EPM)
124 mounted on the end-effector, and a magnetic field sensor array. The robotic arm manipulates
125 the EPM to generate magnetic forces and torques that control the capsule's motion, while the
126 magnetic field sensor array provides continuous 6-DOF pose estimation of the WCE, enabling
127 precise spatial positioning (Figure 1a).

128 The perception module captures images at 1280×720 pixels and 30 frames per second, and
129 obtains 6-DOF pose estimates of WCE through AMagPoseNet³⁷, which achieves a translational
130 error of 1.87 ± 1.14 mm, a rotational error of $1.89 \pm 0.81^\circ$, and a latency of 2.08 ± 0.02 ms.
131 Visual frames and magnetic localization data serve as feedback signals for navigation and
132 control. Within this module, an offline edge-contour-depth fusion pipeline integrates three core
133 components—edge detection, contour matching, and a monocular depth estimation neural
134 network—to extract geometric and structural features that enhance environmental perception
135 and support accurate positioning in complex anatomical environments.

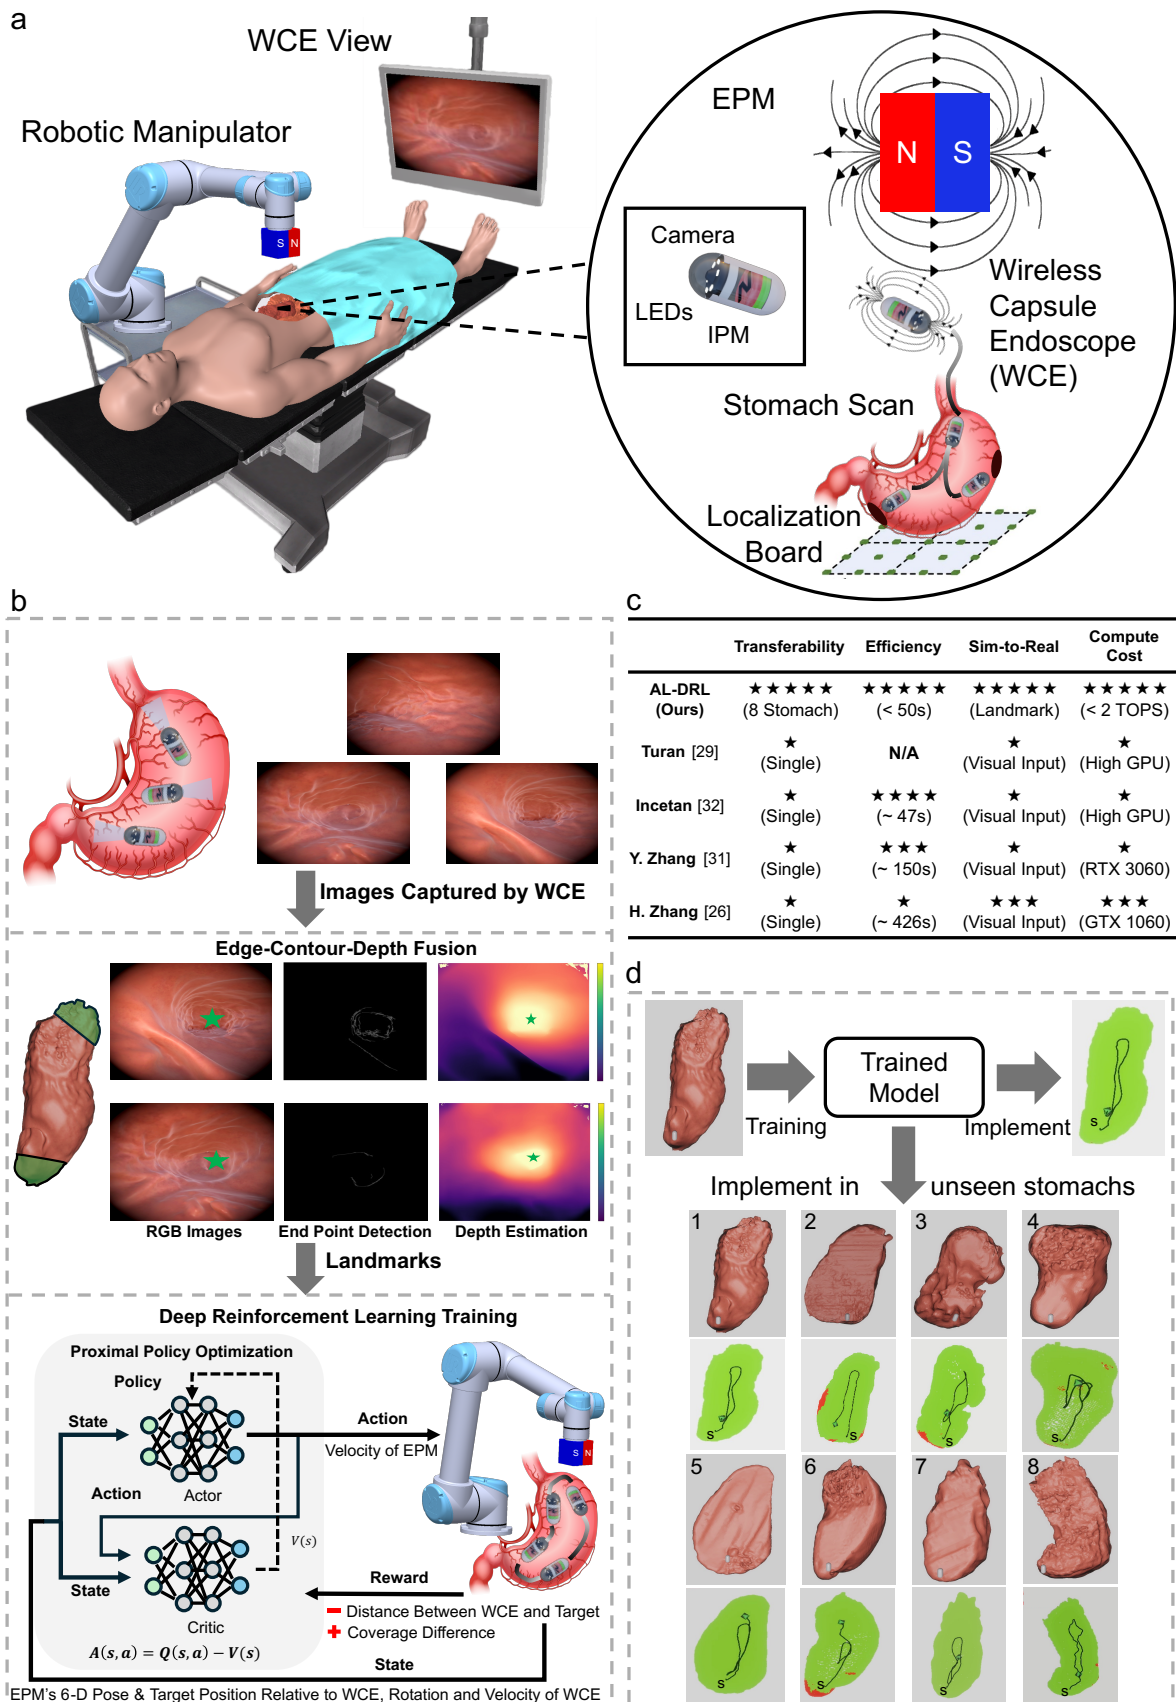
136 The decision-making module is built on a DRL training architecture in which the navigation
137 policy π , parameterized by a neural network, is iteratively refined through agent-environment
138 interaction to maximize the coverage ratio and minimize the elapsed time. A key feature of this
139 architecture is that the DRL policy π does not take high-dimensional image data as input.
140 Instead, the policy's observation vector is composed exclusively of low-dimensional state
141 information, including the 3D coordinates of the target landmark (derived from the perception
142 module) and the EPM and capsule's real-time 6-DOF poses (Figure 1b). This design drastically
143 reduces the computational load and decouples the control policy from visual domain shifts. At
144 each step, the policy outputs EPM displacement and rotation commands, which directly
145 influence the capsule's trajectory. Training is conducted in a simulated stomach reconstructed
146 from patient CT data to preserve anatomical fidelity (Note S1, Supporting Information).

147 In the simulation environment, the AL-DRL framework achieves autonomous, high-coverage
148 navigation across eight distinct patient-derived stomach models (Figure 1d) through the
149 synergistic integration of the perception, decision, and actuation modules. The process begins
150 with the perception module, which fuses real-time 6-DOF pose data with visual priors. These
151 priors are derived from a preprocessing stage that identifies stable, representative gastric
152 landmarks according to established selection criteria—universality, distinctiveness, and
153 navigational utility. These landmarks establish a semi-deterministic global trajectory,
154 embedding anatomical prior knowledge to structure the complex gastric lumen and guide the
155 agent's exploration. The decision module leverages this prior knowledge to guide policy
156 learning. The agent is trained to systematically traverse the landmarks while locally
157 maximizing mucosal coverage, a behavior explicitly shaped by the reward function:

$$r_t = k_1(C_t - C_{t-1}) - k_2D_t + R \quad (1)$$

158 This function incentivizes increases in the coverage ratio C_t via the term $k_1(C_t - C_{t-1})$,
159 penalizes Euclidean distance D_t from the current target landmark, and provides a discrete
160 positive reward R upon reaching a designated endpoint. Upon reaching a landmark, the agent
161 reorients toward the next, decomposing the complex global coverage task into a sequence of
162 manageable navigation sub-goals and ensuring systematic inspection of all gastric regions
163 (Figure 4). The actuation module translates the policy’s commands into precise magnetic forces
164 and torques via the EPM to control the capsule’s motion.

165 By augmenting the efficiency of point cloud-based training with pre-computed landmark and
166 depth features derived from classical image processing and lightweight neural inference, the
167 framework circumvents the need for large scale datasets and computationally intensive training
168 schedules. These embedded visual priors are the key to enhancing transferability, enabling
169 policies trained on a single stomach to be transferred directly to anatomically distinct, unseen
170 stomach models while maintaining high performance.



171
172
173

Figure 1. Experimental system and workflow for transferable DRL-based endoscopic navigation. (a) Schematic of the experimental setup with a single camera capsule inside the stomach and an external 6-DOF robotic arm actuated via a permanent magnet. (b) AL-DRL training pipeline comprising

174 image acquisition, preprocessing, and policy learning via PPO. **(c)** Comparison with State-of-the-Art
175 learning-based gastric navigation methods. **(d)** Training and transferability paradigm: DRL agents are
176 trained with landmark guidance on a single stomach phantom and subsequently tested in previously
177 unseen anatomies to assess model transferability; *s* denoted the start point of the trajectory.

178 **Transferability Characterization of Vanilla DRL Algorithms in WCE Navigation**

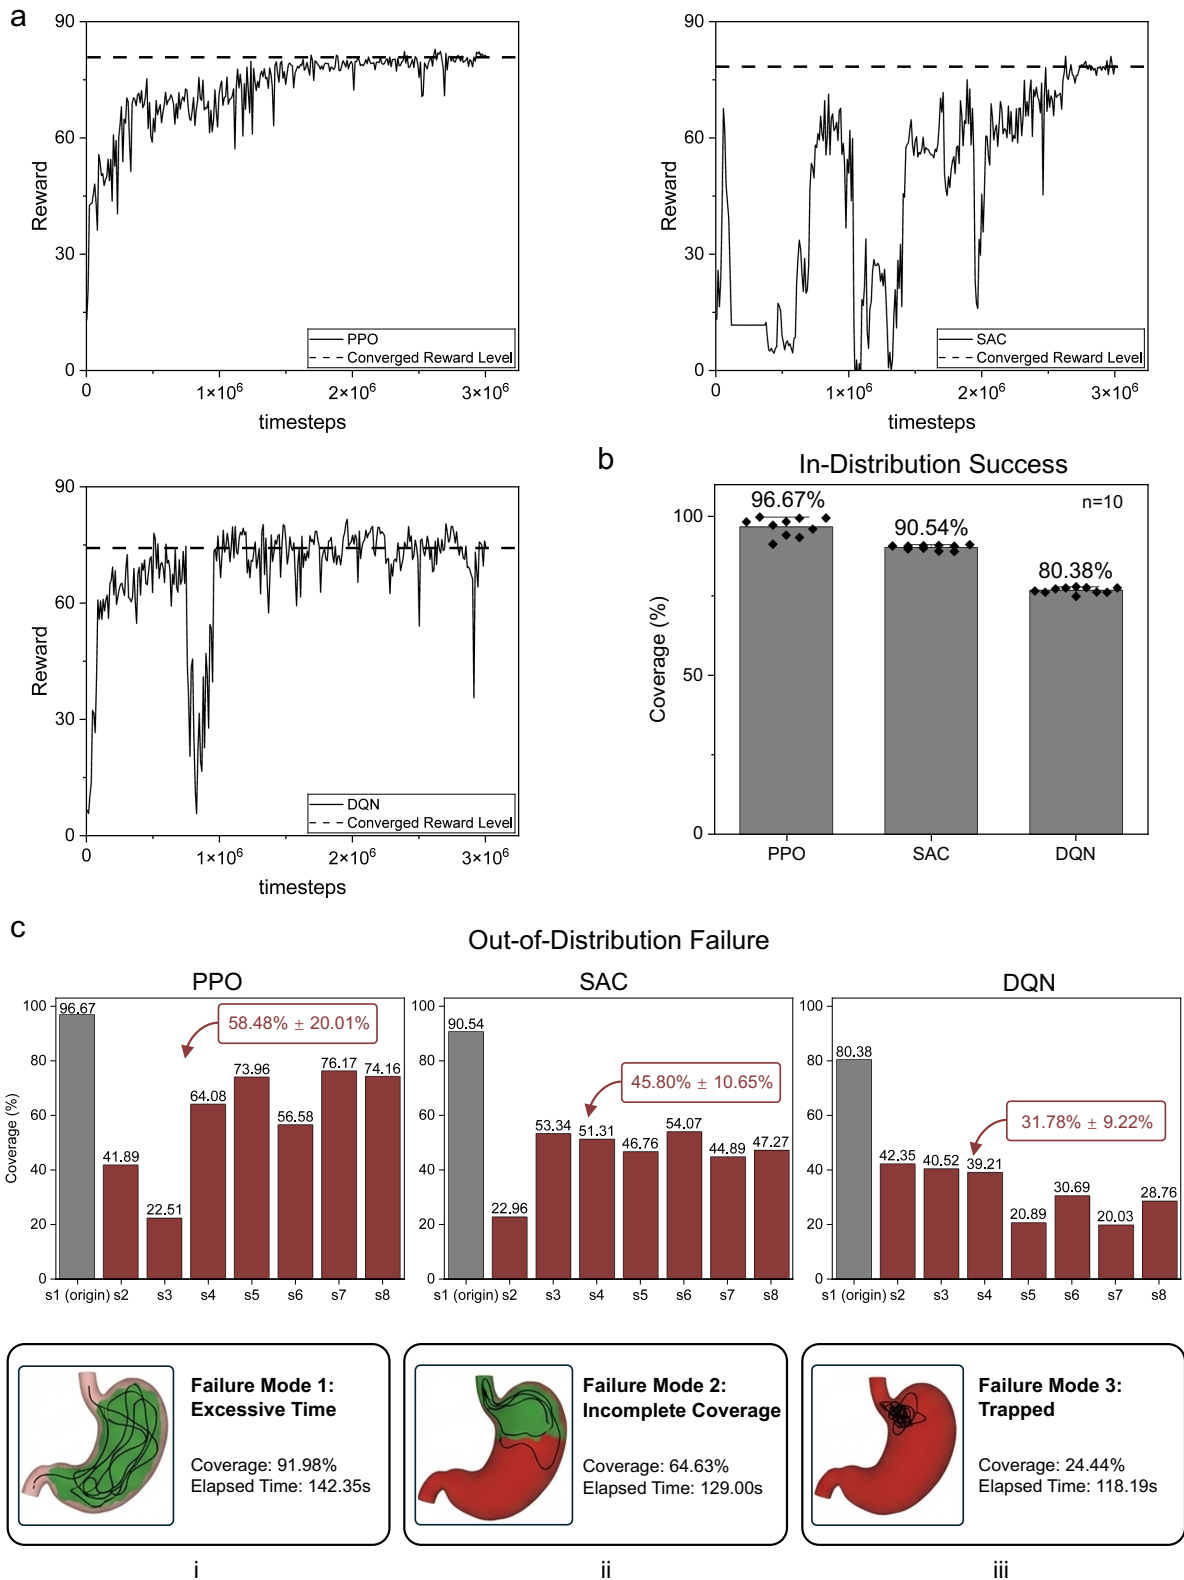
179 To evaluate the transferability capacity of DRL algorithms for gastrointestinal navigation, we
180 benchmarked three vanilla representative state-of-the-art methods—Proximal Policy
181 Optimization (PPO)³⁸, Deep Q-Network (DQN)³⁹, and Soft Actor-Critic (SAC)⁴⁰—within an
182 anatomically realistic stomach simulation environment. PPO, widely recognized for its
183 robustness in complex continuous control tasks^{41,42}, demonstrated superior stability and
184 learning efficiency during training on a single anatomical model. It achieved a mean reward of
185 72.52 with standard deviation (SD) of 10.20, outperforming DQN (mean reward 67.56, SD
186 14.36) and SAC (mean reward 46.53, SD 25.54) under identical training conditions (Figure 2a).
187 This superiority was further validated in deployment. When the trained policies were executed
188 for ten consecutive trials within the same simulation environment used for training (*s*₁), all
189 three DRL agents demonstrated consistently high mucosal coverage with relatively small
190 standard deviations, indicating that the learned policies were stable and effective within the
191 familiar setting. Specifically, the PPO-based agent achieved a mean mucosal coverage of 96.67%
192 (SD 2.93%), which substantially outperformed the agents trained with SAC (90.54% coverage,
193 SD 3.75%) and DQN (80.38% coverage, SD 6.01%). These quantitative data confirmed PPO
194 as the most robust and proficient algorithm for single anatomy navigation, establishing it as the
195 baseline for subsequent transferability studies (Figure 2b).

196 However, this high level of proficiency and consistency proved to be environment-specific.
197 When the trained PPO policy was applied to previously unseen stomach anatomies (*s*₂–*s*₈), the
198 agent’s performance exhibited a dual degradation: a significant reduction in mean mucosal
199 coverage coupled with a dramatic escalation in trial-to-trial variability. Specifically, the mean
200 coverage plunged from 96.67% in the training environment to just 58.48% in unseen
201 environments, while the standard deviation surged from a 2.93% to over 20%. This
202 simultaneous collapse in both efficacy and reliability underscores a profound lack of
203 transferability, indicating that the learned policy is highly sensitive to anatomical variations.
204 Similar patterns of performance decay and increased volatility were observed for SAC (45.80%
205 ± 10.65%) and DQN (31.78% ± 9.22%) (Figure 2c; Note S2, Supporting Information), further
206 confirming that while these vanilla DRL agents can master a specific environment, they fail to
207 maintain a stable and effective navigation strategy across heterogeneous patient structures.

208 A detailed analysis of these cross-anatomy trials revealed distinct and recurring failure modes.
209 For instance, while one agent achieved 91.98% surface coverage, the navigation time increased
210 significantly to 142.35 *s* (Figure 2c(i)). In other cases, seemingly high area scanning masked
211 poor exploration, with the capsule covering mainly the upper stomach while leaving other
212 regions unobserved, resulting in only 64.63% total coverage (Figure 2c(ii)). Critically, agents
213 could become trapped in localized areas, failing to identify new rewarding trajectories and
214 stagnating at just 24.44% coverage after 118.19 *s* (Figure 2c(iii)). Our analysis indicates these
215 failures stem from the policy’s over-reliance on the specific geometric features of the training

216 environment; even small changes in the internal spatial distribution of point cloud coordinates
217 were sufficient to impair navigation performance. These limitations reveal a fundamental
218 deficiency in exploratory flexibility, hindering cumulative reward acquisition and navigation
219 efficiency in unseen environments. In a clinical context, the aforementioned transferability
220 failures are untenable, as reliable detection of pathologies—including ulcers, polyps, and
221 tumors—is predicated on comprehensive and evenly distributed mucosal visualization, while
222 efficient navigation is essential for minimizing examination time and patient discomfort⁴³⁻⁴⁵.

223 These findings expose a fundamental weakness: policies trained solely on single data
224 modalities, such as 3D point clouds, invariably overfit to the specific anatomical structures of
225 the training environment, leading to a critical lack of adaptability in diverse clinical settings.
226 We therefore posit that achieving robust, transferable autonomous navigation requires moving
227 beyond purely stochastic exploration to a paradigm that integrates prior anatomical knowledge
228 or domain specific biomedical cues—such as geometric or landmark features—into the
229 navigation process. Such structured priors can guide the agent along semi-deterministic yet
230 adaptable trajectories, enhancing both coverage efficiency and trajectory stability while
231 preserving transferability across previously unseen stomach anatomies⁴⁶. Achieving this
232 requires the use of visual features that are not only distinctive for reliable detection but also
233 ubiquitous enough to serve as consistent navigational references, enabling the capsule to
234 efficiently and systematically explore the gastric environment.



235 **Figure 2. Training performance and transferability of DRL agents.** (a) Learning curves of PPO,
 236 SAC, and DQN agents during the training phase on a single stomach model. (b) In-distribution success
 237 rates, demonstrating that all three algorithms achieve high coverage with small variance across 10
 238 independent trials when deployed in the original training environment. (c) Out-of-distribution
 239 performance and failure mode analysis in unseen anatomies. The bar charts compare the mean
 240 coverage across eight anatomical scenarios (s1–s8), where s1 is the original environment and s2–s8
 241 are unseen structures. The red labels indicate the aggregate performance across the unseen structures

242 (s2–s8), with data presented as mean \pm SD. The bottom panels visualize three representative failure
243 modes: (i) Excessive Time, (ii) Incomplete Coverage, and (iii) Trapped in limited regions.

244 **Landmark Involvement and Selection Principle**

245 In gastric environments, most of the mucosal surface lacks consistent or distinctive visual
246 markers, which complicates automated navigation. Nonetheless, certain anatomical landmarks
247 exhibit both visual distinctiveness and structural consistency, making them reliable references
248 for guiding an autonomous capsule through complex intraluminal environments. In this study,
249 four landmark pairs—that is, eight gastric landmarks—were examined (Figure 3a). Features
250 such as the fundus, pyloric antrum, cardia, and pylorus possess pronounced edges and
251 well-defined contours, enabling consistent detection by image processing algorithms. By
252 contrast, structures such as the greater and lesser curvature, the angle of His, and the pyloric
253 canal often have irregular or poorly defined borders, reducing their reliability when transferring
254 navigation policies between different anatomies. Selecting prominent and dependable
255 landmarks is, therefore, essential to enhance an agent’s capacity to transfer across anatomical
256 variability.

257 We formalize this selection process through three essential criteria: universality, ensuring
258 landmarks are conserved across patient anatomies; distinction, requiring landmarks to be
259 perceptually salient for reliable detection; and navigation utility, demanding that they anchor
260 efficient, coverage-maximizing trajectories. An evaluation of four candidate landmark pairs
261 across eight patient-derived stomach models was performed to identify the optimal set based
262 on these criteria.

263 The criteria of universality and distinction address perceptual reliability. Our analysis revealed
264 that the fundus, pyloric antrum, cardia, and pylorus possess pronounced edges and well-defined
265 contours, enabling robust identification by image processing algorithms across all models.
266 Quantitative analysis confirmed that the fundus and pyloric antrum, in particular, were detected
267 with 100% accuracy and exhibited maximal edge intensities. In contrast, features such as the
268 greater and lesser curvatures or the angle of His, while anatomically present, lack sharp,
269 consistent borders, rendering them unreliable for policy transfer between anatomies.

270 Perceptual reliability alone, however, is insufficient. The criterion of navigation utility
271 addresses the functional role of landmarks in guiding exploration. For instance, while the cardia
272 and pylorus are morphologically stable, their spatial locations are suboptimal for anchoring
273 trajectories along the primary gastric axis. Their limited visibility from large portions of the
274 lumen makes them ineffective as global endpoints for systematic inspection. Conversely, the
275 fundus and pyloric antrum are situated at opposing termini of the stomach, providing
276 consistently visible endpoints that define a natural axis for full lumen traversal and efficient
277 path planning.

278 Based on this multi-criteria analysis, the fundus and pyloric antrum were identified as the
279 optimal landmark pair, as they uniquely satisfy all three principles of universality, distinction,
280 and navigation utility. This conclusion is quantitatively substantiated by our evaluation, which
281 demonstrated that the fundus and pyloric antrum were detected with 100% accuracy and
282 exhibited maximal edge intensities across all eight patient-derived stomach models (Figure 3b).
283 Their empirically confirmed perceptual reliability, combined with their strategic anatomical
284 placement for guiding a global traversal of the stomach, solidifies their selection. By

285 embedding this principled landmark selection into the DRL framework, we provide the agent
286 with a robust anatomical map, overcoming a critical limitation of methods reliant on featureless
287 geometric data and establishing a foundation for transferable navigation (Note S3, Supporting
288 Information). Accordingly, all subsequent implementations and evaluations of the AL-DRL
289 framework exclusively utilize the fundus and pyloric antrum as the designated anatomical
290 landmarks.

291 **Transferability Characterization of the Offline AL-DRL Framework**

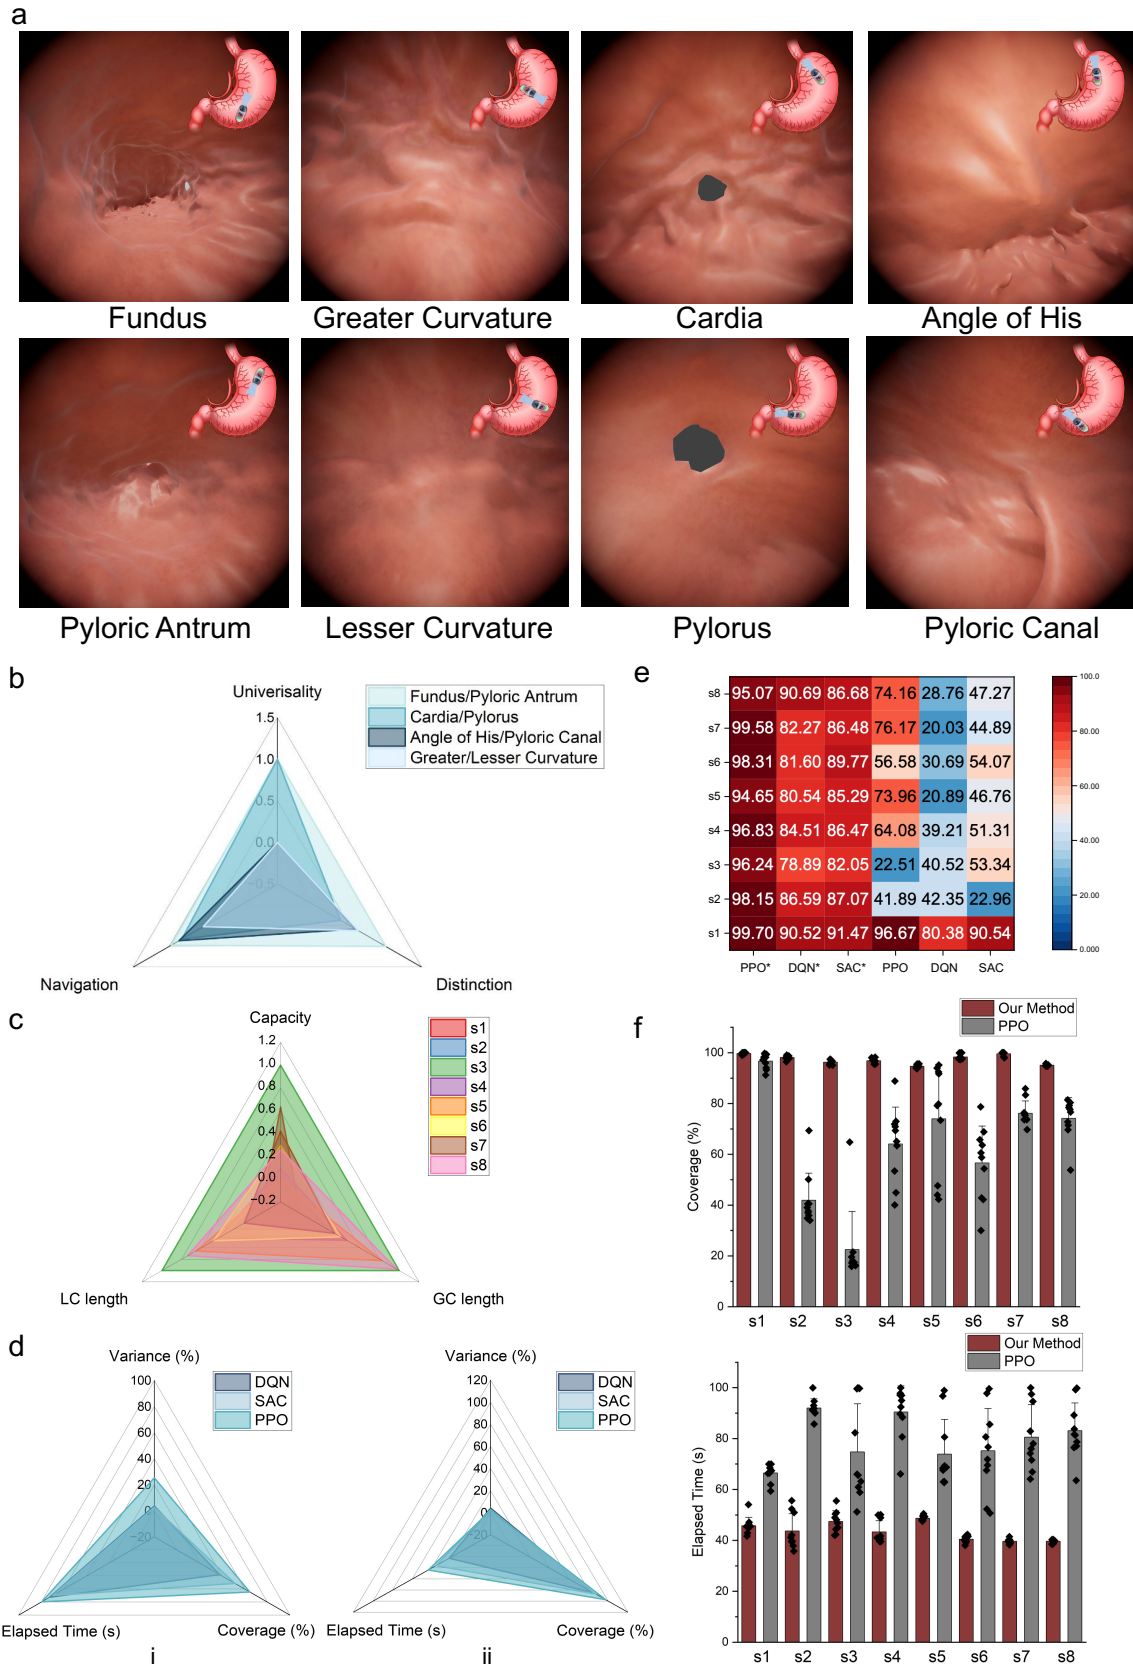
292 Having established a principled strategy for landmark selection, the next challenge is to
293 implement a perception module that reliably identifies these landmarks in real time under the
294 severe hardware constraints of a clinical WCE platform. Unlike domains with rich resources,
295 such as autonomous driving, the WCE operates with a low-power ARM processor (480 MHz),
296 a limited energy budget ($< 3 \text{ W}\cdot\text{h}$), and a strict closed-loop latency requirement ($< 50 \text{ ms}$),
297 rendering computationally intensive deep learning architectures infeasible. To meet these
298 constraints, we developed a computationally efficient edge-contour-depth fusion module. This
299 module combines classical, low-overhead methods—Canny edge detection and Hu moment
300 invariants for rapid landmark identification (10 ms per image pair) (see the “Landmark
301 Identification” section in Methods)—with a pre-trained, lightweight deep neural network
302 (DispNet, 14.84 M parameters) for accurate monocular depth estimation (RMSE of 0.37 cm)
303 (see the “3D Landmark Localization” section in Methods). This approach proved highly robust,
304 achieving 100% landmark detection reliability across all eight distinct patient-derived stomach
305 anatomies. This module thus provides the necessary low-dimensional landmark coordinates for
306 the DRL policy, ensuring the entire navigation framework—from perception to decision-
307 making—operates efficiently within the WCE’s severe computational constraints.

308 To quantitatively validate our framework, we established a robust and clinically relevant
309 experimental testbed. The evaluation was conducted across eight anatomically diverse gastric
310 models reconstructed from patient CT scans, ensuring the simulation captured realistic
311 variations in gastric size, morphology, and mucosal texture (Table S2 and Note S4, Supporting
312 Information). All experiments were performed within a high-fidelity Unity3D simulator that
313 accurately models the physical interactions between the EPM and the WCE. The simulation
314 environment maintained consistent parameters across all trials, including a monocular camera
315 with a fixed 102° field of view and an initial capsule position 30 cm below the EPM to
316 standardize magnetic coupling. The virtual stomach was positioned so that the capsule initially
317 rested against the gastric wall, replicating gravitational and peristaltic settling observed *in vivo*.
318 The EPM was manipulated via a UR5 robotic arm using inverse kinematics, with real-time
319 velocity adjustment to achieve smooth, adaptive capsule control. This comprehensive approach
320 ensured that the evaluation conditions closely mirrored those encountered in practice,
321 providing a reliable basis for assessing robustness and translational applicability.

322 To assess transferability, seven additional models were introduced, each differing in shape,
323 volume, and curvature parameters (Figure 3c). The simulation environment maintained
324 constant camera settings, robotic arm configurations, and WCE control parameters to ensure
325 comparability across trials. The monocular camera retained a fixed 102° field of view (FOV),

326 and the capsule was initialized 30 cm below the EPM at the start of each run to standardize
327 magnetic coupling and avoid instability from extreme proximity or separation. The virtual
328 stomach was positioned so that the capsule initially rested against the gastric wall, replicating
329 gravitational and peristaltic settling observed in vivo. The EPM was manipulated via a robotic
330 arm using inverse kinematics, with real-time velocity adjustment to achieve smooth, adaptive
331 capsule control.

332 Our experimental protocol was designed to rigorously assess transferability. For each of the
333 eight anatomical models, we conducted ten independent trials comparing the performance of
334 standard DRL policies (vanilla PPO, SAC, and DQN) against their counterparts trained using
335 our AL-DRL framework. The results demonstrate a profound improvement in performance and
336 transferability. When guided by the AL-DRL framework, the PPO policy's mean coverage
337 across all eight stomachs reached 97.3% with a low SD of 2.0%; this represents a 53.8%
338 increase in coverage and a 45.3% reduction in navigation time compared to its vanilla
339 counterpart (Figure 3f). The benefits extended across all tested algorithms: the SAC policy's
340 coverage increased by 69.1% (to 86.9%) with a 48.9% time reduction, and the DQN policy,
341 previously ineffective for this task, was transformed with a 123.1% increase in coverage (to
342 84.45%) and a 69.2% time reduction (Figure 3d,e).



343

344

345

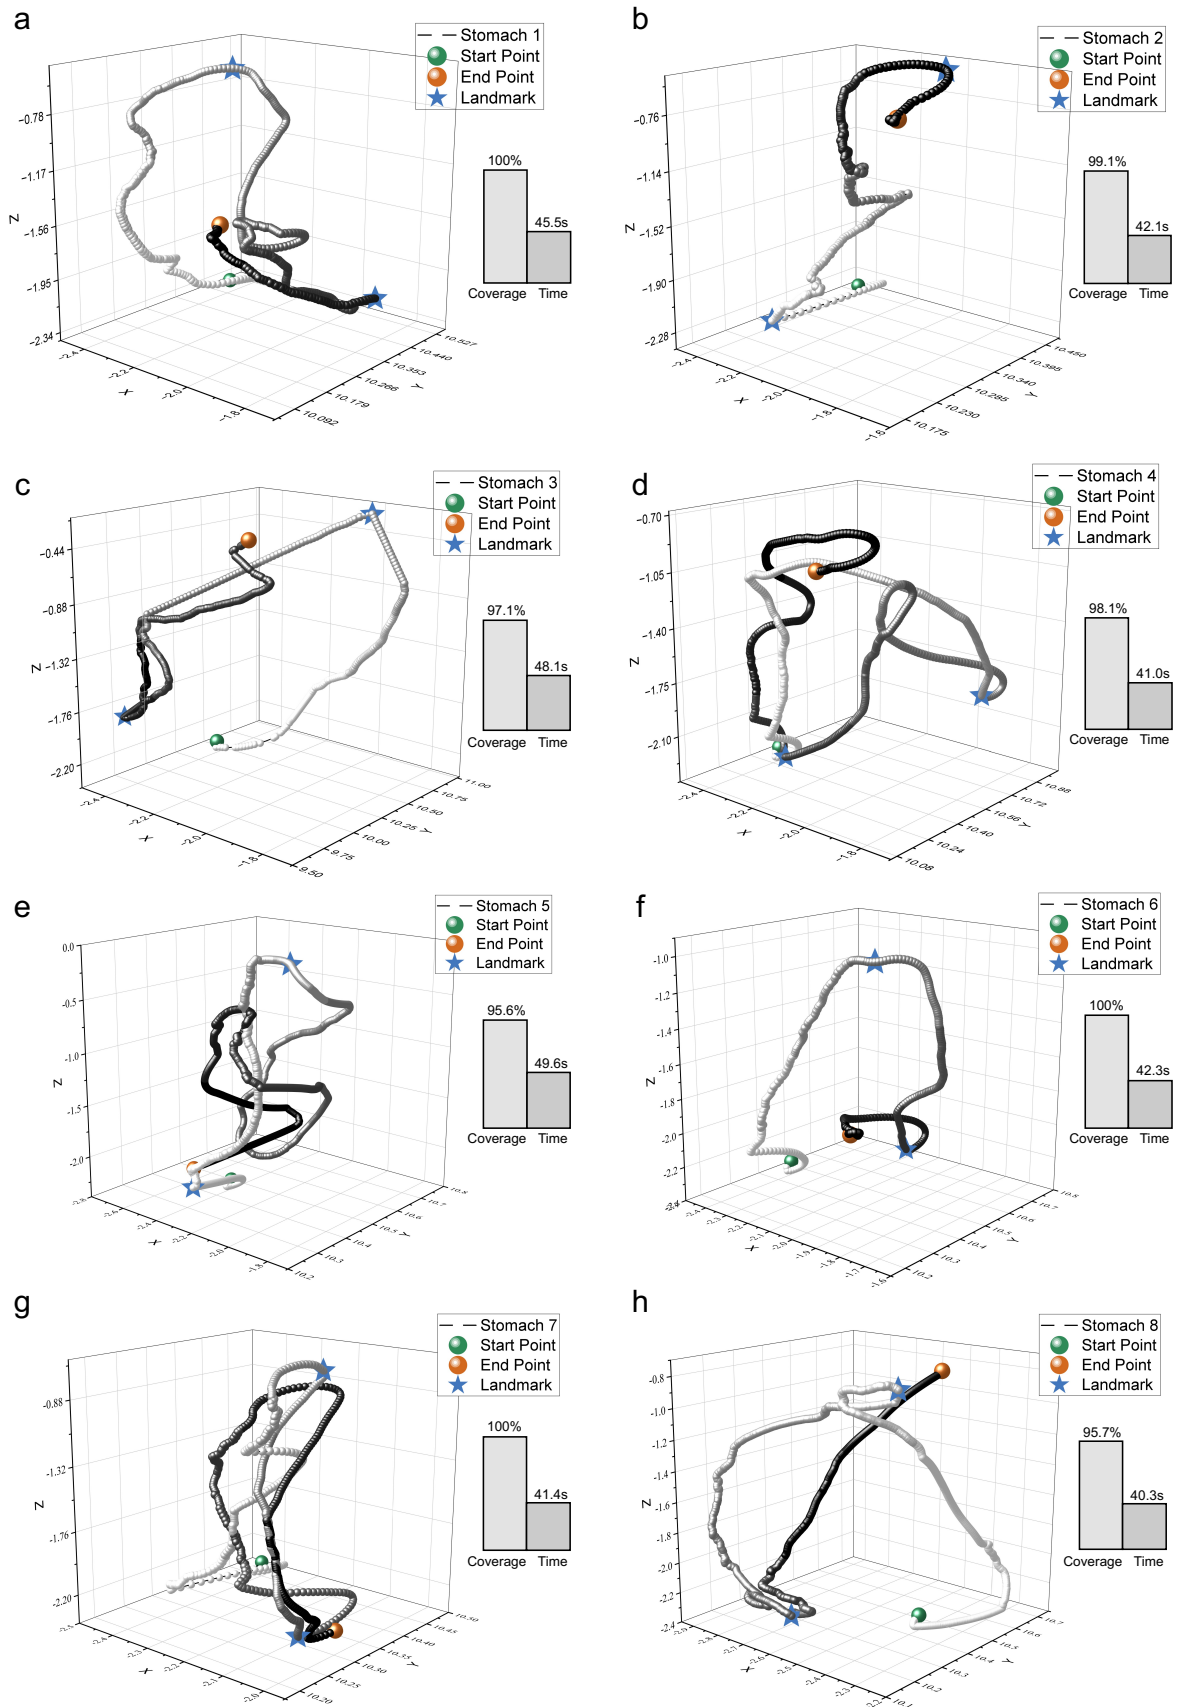
346

Figure 3. Anatomical landmark selection and transferability characterization of AL-DRL on WCE navigation. (a) Intra-gastric views showing candidate landmarks across stomach anatomies. (b) Quantitative evaluation of landmark candidate pairs across three criteria: universality, distinction, and

347 navigation utility. **(c)** Anatomical diversity of patient-derived stomach phantoms: eight models
348 reconstructed from CT data, showing variations in capacity, greater curvature (GC) length, and lesser
349 curvature (LC) length. **(d)** Radar plots comparing key navigation metrics for DRL agents trained without
350 (i) and with (ii) visual landmark guidance across eight distinct stomach models. **(e)** Heatmap comparing
351 the mean gastric coverage of vanilla versus AL-DRL-guided policies across the eight distinct stomach
352 models. **(f)** Coverage ratio and elapsed time variability across eight stomach models for vanilla PPO
353 versus AL-DRL.

354 The success of this landmark-guided approach is visually demonstrated in Figure 4, which
355 shows the 3D navigation trajectories of WCE under the AL-DRL-guided PPO policy, serving
356 as a representative result across all eight anatomically distinct models. The figure reveals a
357 consistent, systematic strategy learned by the agent: upon reaching one anatomical landmark,
358 the policy reliably reorients the capsule toward the next. This behavior effectively decomposes
359 the complex global coverage task into a sequence of manageable, landmark-to-landmark sub-
360 goals. By structuring the exploration process, this strategy prevents the agent from becoming
361 trapped in local regions—a common failure mode for vanilla policies—and ensures a
362 systematic inspection of all major gastric regions. The outcome is rapid and comprehensive
363 coverage, with the agent achieving over 95% mucosal inspection in under 50 s across all eight
364 anatomies.

365 The framework’s superior learning efficiency was also evident in training metrics, where the
366 AL-DRL policy attained a final log-scaled reward of 8.39, overwhelmingly surpassing the 3.26
367 achieved by vanilla PPO (Figure S25, Supporting Information). These findings, supported by
368 both quantitative performance metrics and direct trajectory visualization, confirm that
369 integrating our lightweight visual perception module with anatomical priors systematically
370 resolves the critical issue of policy overfitting, enabling consistent, rapid, and high coverage
371 autonomous navigation across diverse, clinically realistic environments.



372

373 **Figure 4. Robust transferability of anatomical landmark-guided DRL navigation across diverse**
 374 **anatomies.** Superimposed trajectories of the WCE in all eight patient-derived stomach models.
 375 Trajectories generated by the AL-DRL policy with anatomical landmark guidance consistently achieve

376 over 95% gastric coverage within 50 s in each model, underscoring the robustness, efficiency, and
377 reproducibility enabled by visual landmark-based navigation.

378 Hybrid Control for Robust Sim-to-Real WCE Navigation

379 Although the AL-DRL framework performs strongly in simulated gastric environments, direct
380 deployment in real-world WCE navigation is hindered by sim-to-real discrepancies. Our AL-
381 DRL framework is intrinsically robust to these visual discrepancies, as the navigation policy
382 itself is independent of real-time mucosal texture and illumination. The primary sim-to-real
383 gap we must address is therefore related to hardware dynamics and localization, not visual
384 appearance. These dynamic challenges include idealised zero-latency actuation in simulation,
385 which contrasts with physical delays of ~ 30 ms (motor: 10–20 ms; ROS: 8–10 ms), which
386 increase with payload and mechanical friction. Additional issues include kinematic
387 singularities causing >5 mm positional or $>2^\circ$ orientational errors, and magnetic localisation
388 drift (>7 mm, $>5^\circ$) from EPM electromagnetic interference. As the DRL model tolerates $<5\%$
389 input noise, such deviations can trigger rapid performance loss and unstable behavior.
390 Therefore, a direct, closed-loop sim-to-real deployment of the DRL policy is infeasible. To
391 address these challenges, we propose a two-stage control architecture:

- 392 (1) Patient-specific digital twin: Reconstructed from CT data to capture gastric geometry and
393 lumen topology. Physical parameters—FOV 102° , focal length 0.57 mm, illumination 1 lux,
394 magnetic moments (EPM: 119.36 Am^2 , IPM: 0.55 Am^2), reflectivity 0.4, gravity—are
395 empirically calibrated. Realistic collision and friction are modelled in Unity + PhysX using
396 measured Young’s modulus and local wall friction. CycleGAN-based style transfer aligns
397 simulated textures with in vivo imagery.
- 398 (2) Hybrid execution: Policies are trained in simulation and deployed via offline trajectory
399 planning + online adaptive correction. An optimal path is computed in the digital twin, then
400 followed by a model-free ADP controller with real-time visual feedback for deviation
401 correction. Offline planning reduces execution frequency from 200 Hz to 50 Hz, while
402 trajectory smoothing (time/acceleration/jerk minimization) avoids singularities. This
403 confines heavy optimisation to simulation, enabling embedded execution on < 2 TOPS
404 processors.

405 In the second stage of the framework, motion planning results generated in simulation are
406 transferred to the physical system via an adaptive dynamic programming (ADP)-based robust
407 tracking controller that compensates for environmental uncertainties and proactively
408 neutralizes hardware latency through forward state prediction (Note S13, Supporting
409 Information). Unlike conventional proportional-integral-derivative (PID) or fuzzy logic
410 controllers, ADP employs an iterative value-function approximation to optimize control inputs
411 without requiring explicit system dynamics⁴⁷. This model-free property makes it particularly
412 suitable for WCE navigation, where complex actuation dynamics and sensor noise limit the
413 applicability of model-based approaches. The controller minimizes the cost function

$$V = \int_t^\infty (Q(x) + u^T R u) d\tau, Q(x) > 0, R > 0 \quad (2)$$

414 where $x = [p_x - p_{xd}, p_y - p_{yd}, p_z - p_{zd}, \phi - \phi_d, \theta - \theta_d, \psi - \psi_d]^T$ represents the position
415 and rotation tracking error state, and $u = [v, \omega]^T$ denotes the linear and angular velocity

416 control input of the EPM. Here, $Q(x)$ penalizes trajectory deviations, while R regularizes
 417 control effort to ensure smooth actuation.

418 At each iteration j , the control policy $\hat{u}_j(x)$ is updated by solving:

$$\hat{c}_{j+1} = \arg \min_{c \in \mathcal{C}} \left\{ \underbrace{\hat{b}_j^T \rho(x, c^T \sigma(x))}_{\text{Value Gradient}} + \underbrace{(Q(x) + \hat{u}_j(x)^T R \hat{u}_j(x))}_{\text{Cost Penalty}} \right\} \quad (3)$$

419 where \hat{c}_j and \hat{b}_j are neural network weights updated from historical trajectory data via
 420 Recursive Least-Squares (RLS). The term $\hat{b}_j^T \rho(\cdot)$ estimates the gradient of the value function
 421 $V_j(x)$, while $Q(x) + \hat{u}_j(x)^T R \hat{u}_j(x)$ penalizes tracking errors and excessive control effort.
 422 Through successive updates, ADP converges to a policy capable of rejecting disturbances while
 423 maintaining stable trajectory tracking (Figure 5a).

424 Real-time robustness is further enhanced by coupling ADP with online anatomical landmark
 425 detection for dynamic trajectory correction. External disturbances such as magnetic
 426 localization drift, gastric peristalsis, and mechanical actuation latency can cause substantial
 427 path deviations that threaten task completion. Informed by the reward function's emphasis on
 428 endpoint traversal and turning manoeuvres, the system uses real-time endpoint recognition as
 429 a corrective reference. When the deviation from the target exceeds allowable limits, a new path
 430 is computed toward the nearest endpoint, and the controller actively steers the WCE back to
 431 the intended trajectory. Once the capsule re-enters a 5 mm proximity zone around the target,
 432 the system resumes optimal coverage tracking. This dual-layer approach mitigates large-scale
 433 deviations and sustains high coverage accuracy despite environmental perturbations.

434 We validated the ADP framework through comparative experiments against PID ($K_p = 0.5$,
 435 $K_i = 0.1$, $K_d = 0.05$) and fuzzy control strategies. An optimal trajectory trained in simulation
 436 was executed under each method, while robustness was challenged by introducing magnetic
 437 field fluctuations ($\pm 2\%$) and localization noise (positional ± 3 mm, orientation $\pm 2^\circ$). As shown
 438 in Figure 5c, ADP consistently achieved the lowest mean squared tracking error
 439 ($\text{MSE} = 0.04 \text{ cm}^2$) compared with PID (0.57 cm^2) and fuzzy control (0.31 cm^2), reflecting an
 440 89.5% reduction in error relative to PID. Although a slightly higher error was observed in the
 441 y-z plane (Figure 5e), overall accuracy and robustness remained superior. In high-curvature
 442 navigation segments (radius < 12 mm), characteristic of abrupt orientation changes required for
 443 complete mucosal scanning, ADP maintained a maximum deviation of only 1.3 mm,
 444 substantially outperforming PID (5 mm) and fuzzy control (3 mm) (Figure 5b,d).

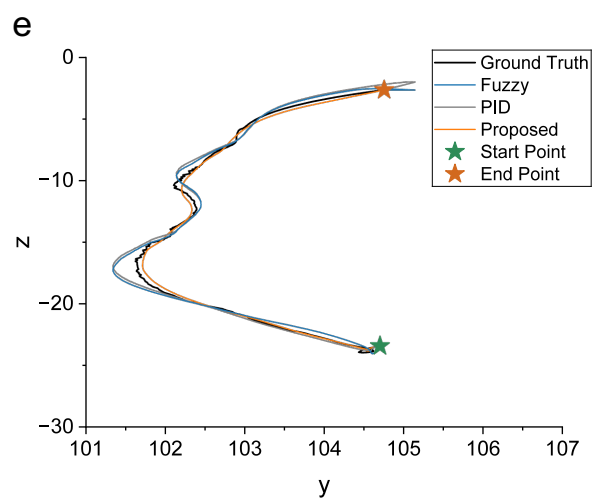
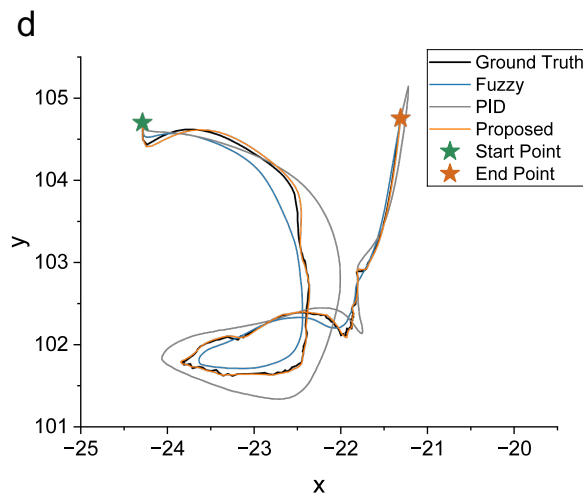
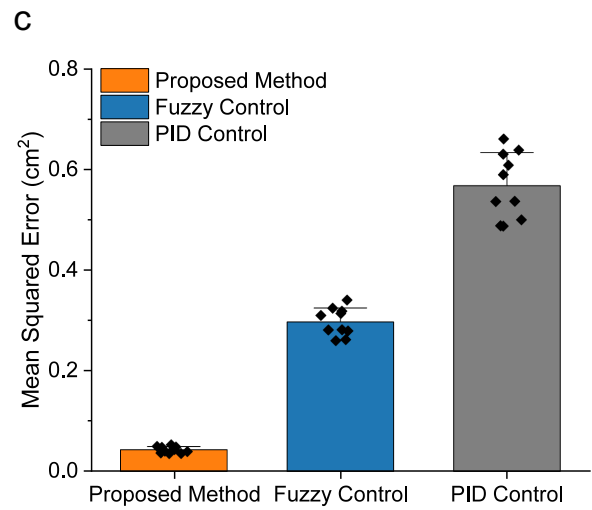
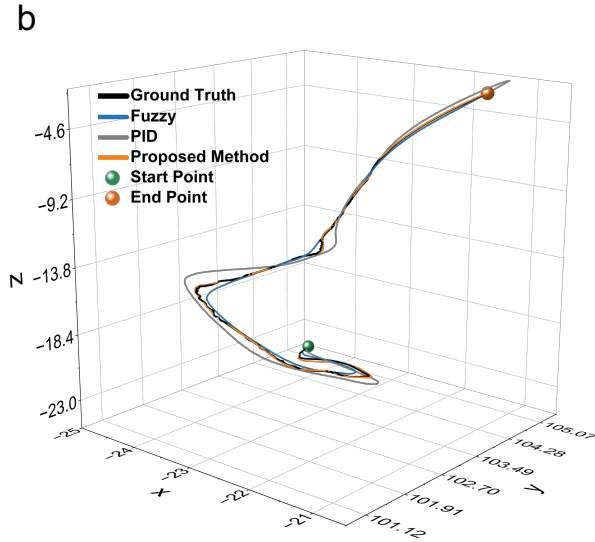
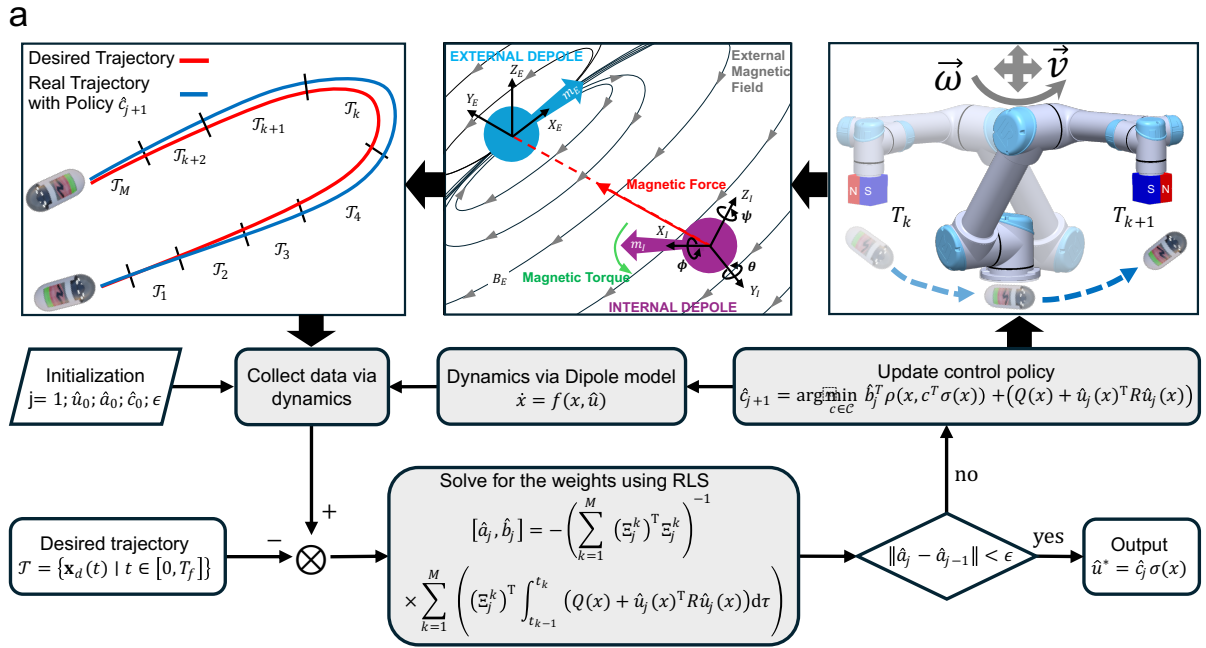


Figure 5. Robust trajectory tracking and control performance under simulated perturbations. (a) Schematic of the offline trajectory planning workflow using a digital twin, coupled with the ADP-based

447 robust tracking controller for real-time correction. (b)-(c) Comparative analysis of mean squared
448 tracking error under simulated perturbations for ADP, PID, and fuzzy control methods. (d)-(e)
449 Orthogonal projections of the 3D trajectory shown in (b) onto the x-y and y-z planes.

450 **Real-World Ex-vivo Validation of the Offline AL-DRL Framework**

451 To validate the proposed framework under clinically representative conditions, we developed
452 a three-dimensional multimodal ex-vivo test platform (Figure 6b and the “Simulation and Real
453 Environment” section in Methods). The complete data-to-execution pipeline is outlined in
454 Figure 6a. Two-dimensional computed tomography (CT) images of the stomach were first
455 acquired and segmented using advanced algorithms to isolate the organ from surrounding
456 tissues. The segmented slices were reconstructed into high-fidelity three-dimensional models
457 using 3D Slicer, and anatomical realism was enhanced by mapping RGB texture information
458 from endoscopic imagery via style-transfer techniques. The resulting patient-specific gastric
459 models were then imported into the simulation environment for agent training (Figure 6c), with
460 simulation outputs subsequently translated into real-world robotic manipulator commands by
461 an adaptive dynamic programming (ADP)-based controller.


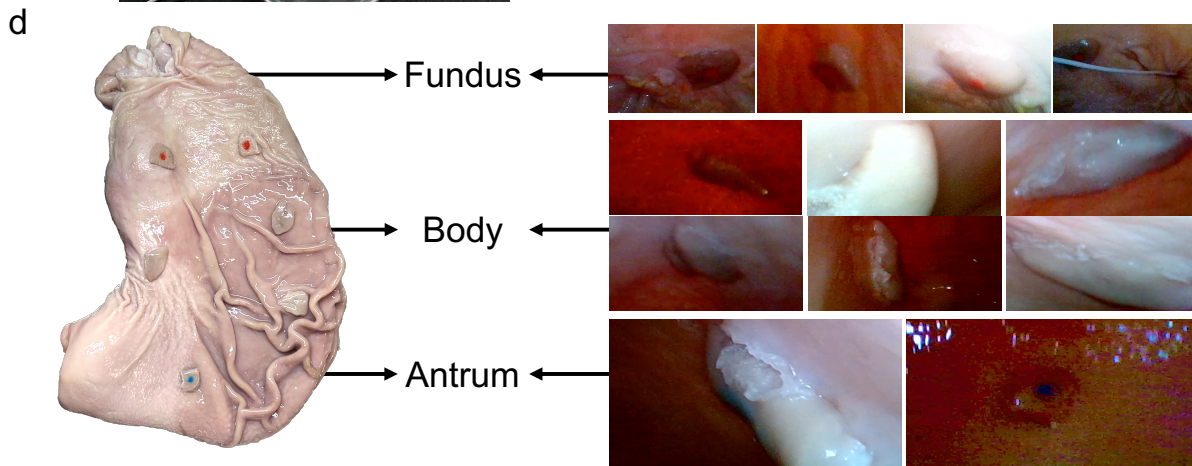
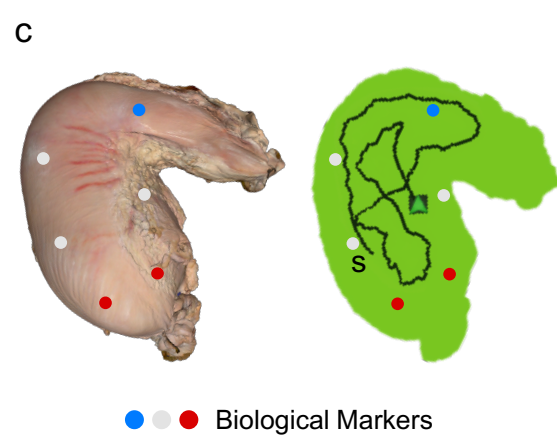

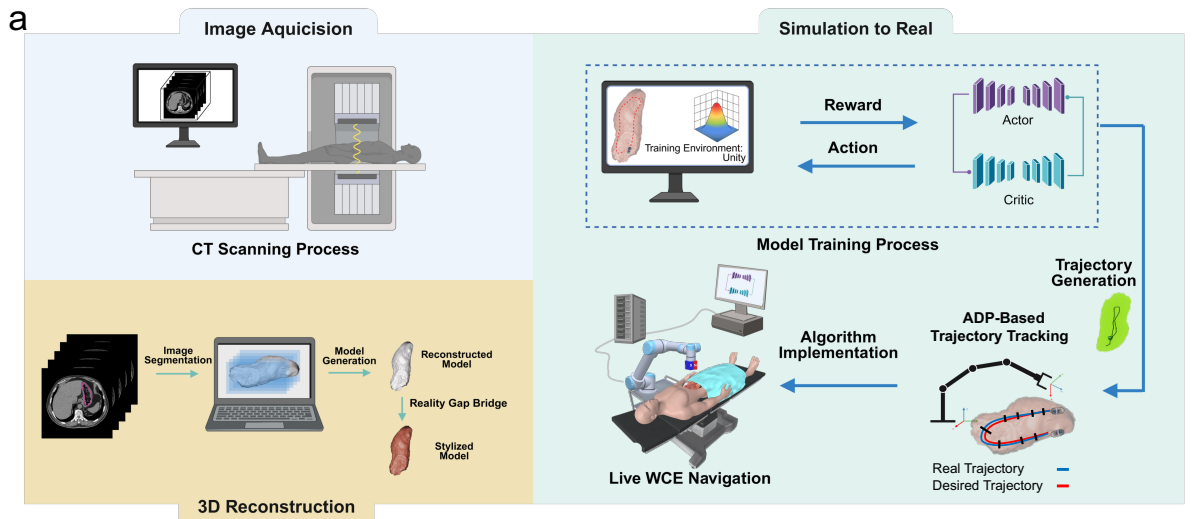
462 To ensure standardized performance assessment, 12 biological markers, crafted from excised
463 gastric tissue to simulate tumors, were strategically distributed throughout the ex-vivo pig
464 stomach. To ensure a comprehensive and uniform assessment, the markers were placed on the
465 surfaces of the three primary gastric regions—the fundus, body, and antrum—to provide fixed
466 references for quantitative coverage and trajectory accuracy. Specifically, four markers painted
467 red were placed in the fundus, six unpainted markers were placed in the body, and two markers
468 painted blue were placed in the antrum. Figure 6d depicts the markers layout on one side; an
469 identical, symmetrical arrangement was placed on the opposing side. The 3D stomach model
470 reconstructed from a scanner is imported into Unity as the training environment. To establish
471 a high-fidelity digital twin of the experimental setup, our simulation platform was developed
472 in the Unity engine. The physical parameters of every component—including dimensions,
473 mass, magnetic properties, and camera intrinsics—were meticulously configured to match their
474 real-world values, thereby minimizing the sim-to-real gap. After AL-DRL training, the
475 optimized trajectory was deployed in reality with the UR5 arm. To achieve the target linear
476 and angular velocities for the EPM, the corresponding joint velocities of the manipulator were
477 calculated via inverse kinematics using the manipulator Jacobian (Note S5, Supporting
478 Information). During this process, a manipulability index was maintained above 0.1 to avoid
479 singularities. Figure 6d shows the anatomical distribution of the markers and the corresponding
480 WCE images acquired during navigation, confirming successful detection in all gastric regions.

481 The scanning protocol (Figure 6e) prescribed a serpentine trajectory designed for
482 comprehensive surface coverage. The path originated in the gastric body, followed the greater
483 curvature to the antrum and pylorus, returned along the lesser curvature toward the fundus, and
484 finally swept back toward the gastric body to complete the inspection. Orientation adjustments
485 maximized mucosal exposure throughout. The first segment along the greater curvature took
486 40 s, covering about one-third of the gastric surface; the second, from the antrum along the
487 lesser curvature to the mid-body, lasted 80 s due to a longer distance and a required 180° turn;

488 the final segment, covering the fundus and remaining body, took 40 s with a similar rotation
489 event despite a shorter path.

490 Reliability and repeatability tests over five independent scans yielded an average coverage ratio
491 of 87% (number of captured markers: 12/12, 10/12, 11/12, 9/12, 11/12) with a completion time
492 of 139.8 ± 10 s (coefficient of variation 6.5%). In comparison, a trained human operator using
493 a joystick required 298.4 ± 30 s (coefficient of variation 7.1%) for equivalent coverage,
494 representing a 53% time reduction with the autonomous system. This improvement stems from
495 the ADP controller's stability, particularly in regions of high curvature (radius < 20 mm), where
496 tracking accuracy remained at 1.5 mm. The singularity-avoidance constraint prevented abrupt
497 robotic arm movements, ensuring smooth actuation. [Collectively, these results demonstrate the
498 potential of the proposed platform as a promising framework for autonomous gastrointestinal
499 navigation. The consistent performance across diverse simulations and initial physical tests
500 suggests a viable path toward reducing operator dependency through intelligent endoscopic
501 robotics.](#)

502



503
504
505

Figure 6. Validation of sim-to-real transfer and autonomous navigation within an ex-vivo pig stomach model. (a) Simulation-to-reality transfer pipeline illustrating the training of a DRL agent in a

506 simulated environment and subsequent deployment in a real-world WCE procedure. **(b)** The
507 experimental setup, comprising a UR5 robotic arm that manipulates an EPM to actuate the tethered
508 modular capsule. **(c)** The ex-vivo stomach model, reconstructed from 3D scan data, and the optimal
509 trajectory generated in simulation and subsequently executed on the physical system via the hybrid
510 control algorithm. **(d)** Representative endoscopic images from the capsule camera, showing markers
511 designating distinct anatomical landmarks during the autonomous scanning process. **(e)** The trajectory
512 and time stamp of the autonomous scan demonstrate high fidelity to the planned path from simulation.

513

514 **Discussion**

515 Here, we present an autonomous WCE area-coverage framework that integrates DRL,
516 anatomical landmark detection, and monocular depth estimation into a cohesive control system.
517 Conventional RL-based navigation approaches often show limited transferability across
518 diverse, patient-specific stomach anatomies, making retraining for each case both
519 computationally demanding and economically impractical in clinical workflows. To address
520 this constraint, we selected PPO—which demonstrated strong single-model navigation
521 performance—and enhanced its transferability by streamlining the training pipeline and
522 embedding visual feature extraction mechanisms that produce stable, anatomically informative
523 representations. The resulting framework consistently achieved more than 97% mucosal
524 coverage within 50 s across eight anatomically distinct stomach models, effectively eliminating
525 the coverage penalty typically observed when transferring RL agents between heterogeneous
526 anatomies. This performance surpasses point-based and geometry-only approaches,
527 establishing a new benchmark for autonomous endoscopic navigation.

528 Central to this capability is our framework’s core architectural design, which deliberately
529 circumvents the computationally demanding and inherently brittle “end-to-end” paradigm
530 common in visual-DRL systems. This conventional approach, which feeds high-dimensional,
531 real-time endoscopic imagery directly into the policy, suffers from two critical translational
532 barriers: (1) it demands substantial computational resources (e.g., high-performance GPUs)
533 that are infeasible for low-cost, real-time clinical deployment, and (2) it is fundamentally
534 vulnerable to domain shifts, failing to transfer when faced with the vast differences in mucosal
535 appearance across patients. Our framework, in stark contrast, severs this dependency. The
536 policy is trained to operate exclusively on low-dimensional 3D coordinates derived from pre-
537 identified anatomical landmarks. Guided by our principled landmark selection criteria
538 (universality, distinction, and navigational utility), this design is the key to overcoming both
539 aforementioned barriers, yielding profound advantages for clinical translation. First, it renders
540 the navigation task computationally tractable, enabling efficient training and, more importantly,
541 high-speed deployment on low-cost, power-constrained hardware. **By optimizing the execution**
542 **frequency at 50 Hz, we balance efficiency with stability: higher rates are redundant given the**
543 **30 ms physical latency, whereas lower rates cannot sufficiently reject real-time disturbances,**
544 **such as magnetic drift or peristalsis, leading to tracking failures in high-curvature segments.**
545 **This ensures that corrective actions maintain the required control accuracy for robust**
546 **navigation while minimizing the computational burden on low-power hardware.** Second, and
547 most critically, it achieves robust transferability by making the policy intrinsically invariant to

548 the visual domain shifts that plague end-to-end methods. The policy’s performance becomes
549 independent of the vast differences in mucosal appearance across patients, ensuring consistent
550 operation across diverse patient anatomies. By decoupling the complex perception challenge
551 from the real-time control policy, our framework achieves superior robustness and scalability.
552 This anchoring to stable anatomical coordinates, rather than transient visual features,
553 drastically reduces the dimensionality of the policy’s observation space, simplifying
554 optimization and ensuring predictable, transferable behavior.

555 There are, however, notable limitations. The current endpoint-matching algorithm is sensitive
556 to variations in illumination and affine transformations, leading to occasional detection errors
557 when landmarks are imaged from extreme viewpoints or under challenging lighting. [While
558 additional reliability tests demonstrate that our perception module consistently identifies
559 landmarks under off-axis tilts and dim lighting, we acknowledge that extreme in-vivo artifacts
560 \(dense mucus coverage or intense specular reflections\) could still degrade real-time detection.](#)
561 [To ensure safety during perception failures, the system employs temporal prediction for
562 transient occlusions via ADP controller \(<500 ms\) and an active search mode via in-place
563 rotation for sustained losses \(>2 s\), with a manual fallback alert as a final safeguard. More
564 resilient computer vision methods that maintain accuracy across a broader range of
565 intraoperative conditions are needed to address this weakness.](#)

566 [Furthermore, while our ex-vivo experiments utilize a highly flexible tether with negligible
567 macroscopic actuation resistance \(~0.04 N·mm\), it is important to acknowledge that the tether
568 intrinsically acts as a stabilizing “tailfin” from a dynamics perspective. It provides mechanical
569 damping that suppresses high-frequency oscillations, stick-slip friction-induced chatter, and
570 micro-tumbling. In a wireless clinical scenario, the posture dynamics of the capsule will be far
571 more complex and highly non-linear, masking the complex 6-DOF tumbling that a free capsule
572 would exhibit. Additionally, while the system employs active compensation for EPM
573 interference, the reliance on an analytical dipole model presents inherent accuracy limitations
574 when encountering non-linear environmental distortions \(e.g., hard-iron/soft-iron effects\) or
575 dynamic noise from the 6-DOF robotic arm.](#)

576 In addition, although the system demonstrates strong transferability across healthy anatomical
577 variants, validation in larger patient cohorts—especially in cases with pathological alterations
578 or atypical morphology—is essential to confirm clinical robustness. [Under extreme
579 pathological conditions like large hiatal hernias, where landmark spatial relationships may be
580 significantly altered, the framework ensures safety through a confidence-based detection
581 mechanism: if the matching similarity score \(\$D\$ \) falls below a reliable threshold, the system
582 triggers an anomaly alert to prevent erroneous navigation. Furthermore, the low-level ADP
583 controller’s mechanical constraints and the possibility of immediate manual intervention
584 provide redundant safety layers to handle such severe morphological distortions. For the
585 fundamental DRL algorithm, the implemented architecture is based on recurrent neural
586 networks \(RNN\), which is ineffective for long-term dependencies and only indirectly accesses
587 the past via a compressed hidden state.](#)

588 Future work will focus on enhancing the policy’s long-term reasoning capabilities and its
589 adaptability to more dynamic clinical scenarios, while preserving the framework’s core
590 computational advantages. Replacing RNN with a Transformer-based architecture—which
591 excels at modelling complex sequences—could significantly improve the agent’s ability to
592 optimize long-horizon trajectories. Critically, such a model would still operate on the low-
593 dimensional sequence of coordinate and pose data, maintaining the low-cost, transferable
594 nature of our approach. We also plan to address the transition to a wireless state as a critical
595 translational boundary. This involves updating our training pipeline with stochastic noise
596 models to simulate wireless instability and enhancing the ADP controller to specifically
597 address high-frequency perturbations and the reduced damping inherent to tether-free boundary
598 conditions. To further mitigate unpredictable environmental magnetic noise, we intend to
599 integrate an advanced anti-interference framework based on MagRobustNet⁴⁸. By utilizing
600 self-supervised anomaly detection and measurement recovery, this framework has
601 demonstrated a 76.2% improvement in position accuracy, ensuring reliable tracking in complex
602 clinical settings.

603 Additionally, we plan to extend the framework from pre-defined landmark traversal to dynamic,
604 target reaching capabilities, enabling the agent to navigate to specific 3D coordinates of a
605 pathology (e.g., a polyp) identified intraoperatively. To enhance clinical universality, we will
606 expand the training pipeline by incorporating a broader range of pathological gastric models
607 featuring severe morphological distortions, such as hiatal hernias or massive intraluminal
608 compressions. This expansion will enable the AL-DRL policy to adapt to atypical anatomical
609 priors where standard spatial relationships are disrupted. Future efforts will develop more
610 robust, adaptive target identification methods and incorporate recovery strategies for
611 navigating in atypical or pathological anatomies where targets may be obscured.

612

613 **Methods**

614 **Simulation and Real Environment**

615 The simulation environment was developed in Unity (version 2022.3.32f1) as a real-time 3D
616 platform, using the built-in NVIDIA PhysX engine to model rigid-body dynamics with both
617 discrete and continuous collision detection. DRL training was supported through integration of
618 the ML-Agents toolkit (release 20). A 6-DOF Universal Robots manipulator served as the
619 simulated actuator mounting an EPM to control the WCE. Eight stomach models derived from
620 patients were imported into Unity as primary test environments. As these virtual stomach
621 models are represented by point clouds, the coverage ratio (C_t)—a key metric for training—is
622 defined as the percentage of the total point cloud points scanned by the capsule’s camera. To
623 ensure clinical relevance, a point is only recorded as “scanned” if its distance from the camera
624 lens is within a 30 mm threshold, which is informed by clinical evidence demonstrating that
625 magnetically controlled capsule endoscopy (MCE) achieves an overall diagnostic accuracy of
626 93.4% within a 0–60 mm viewing range⁴⁹. Physical parameters were meticulously calibrated
627 to mirror real-world conditions: gravitational acceleration was set to 9.81 m/s², the EPM
628 magnetic moment to 119.36 Am², and the IPM magnetic moment to 0.55 Am². Friction

629 coefficients were set to 0.1 for the gastric wall and 0.5 for the capsule surface. Optical
630 properties were configured with a monocular camera field of view of 102° , a focal length of
631 0.57 mm, material reflectivity of 0.4, and ambient illumination of 1 lux. To simulate gastric
632 peristalsis, wall friction coefficients were regionally varied and subjected to random
633 perturbations during runtime (Note S6, Supporting Information).

634 The physical testbed (Figure 6b) comprised a 6-DOF Universal Robots arm manipulating an
635 EPM via a custom-designed 3D-printed container that mounted a cubic NdFeB magnet ($50 \times$
636 50 mm) at the end-effector, providing an operational workspace radius of 850 mm. The custom
637 WCE measured 12 mm in diameter and 35 mm in length, integrating an OV9734 fibre-optic
638 imaging module (1280×720 pixels at 30 fps, 102° FOV) and an IPM formed by four
639 orthogonal NdFeB pole groups. The capsule was enclosed in a 3D-printed clear resin sheath
640 (Young's modulus 4 GPa, friction coefficient 0.5, tensile strength 70 MPa). This modular
641 design, costing under \$50 to assemble, replaced commercial WCE units that cost over \$1,000
642 and lack real-time video. A low-modulus tether (2 mm diameter, Young's modulus 10 MPa)
643 preserved realistic capsule kinematics, with dynamic torsion tests showing an added torque of
644 only $0.04 \text{ N}\cdot\text{mm}$ under $\pm 180^\circ$ rotation—around 10% of the $0.35 \text{ N}\cdot\text{mm}$ torque generated by
645 the EPM. [Further experimental characterization across a 25–125 mm range confirmed that](#)
646 [magnetic torque remains consistently above the tether's resistive threshold, reaching the](#)
647 [standardized working value of \$0.35 \text{ N}\cdot\text{mm}\$ at approximately 90 mm](#) (Note S7, Supporting
648 [Information](#)).

649 The scanning medium was an ex-vivo pig stomach measuring 31.2 cm in length and 14 cm in
650 maximum width. An air pump was used to inflate the stomach to a distended state to ensure a
651 realistic scanning environment. Twelve biological markers, crafted from excised gastric tissue
652 to simulate tumors, were affixed to the internal ex-vivo pig stomach surface—four designated
653 by red marks in the fundus, six unmarked markers in the body, and two designated by blue
654 marks in the antrum. Capsule localization was achieved with a 5×5 magnetometer array
655 (LIS3MDL, STMicroelectronics; $\pm 1200 \mu\text{T}$ range) mounted on a custom localization board
656 with 40 mm sensor spacing. This array resolves the capsule's 6-DOF pose using AMagPoseNet
657 (Note S8, Supporting Information). Data acquisition was handled by an STM32H743
658 microcontroller, with real-time transmission to a host PC. To compensate for magnetic
659 interference from the EPM, its real-time pose—calculated via robotic forward kinematics—is
660 used in a magnetic dipole model to compute its interference field, which is then subtracted
661 from the sensor array's total measurements to isolate the pure signal from the capsule's internal
662 magnet (Note S9, Supporting Information).

663 **Landmark Identification**

664 To enable efficient landmark identification under constrained computational resources and in
665 the context of point cloud-based stomach modelling, the image processing module is designed
666 to identify and match anatomical landmarks using a template-driven approach. The module
667 comprises two main stages: edge detection and feature matching.

668 For edge detection, we implement the Canny method, a gradient-based algorithm for
 669 delineating object boundaries. Given a grayscale input image $I(x, y)$, horizontal and vertical
 670 gradients are computed via Sobel filters as

$$G_x = \frac{\partial I}{\partial x}, \quad G_y = \frac{\partial I}{\partial y} \quad (4)$$

671 where G_x and G_y are the gradients along the x and y axes respectively. The gradient
 672 magnitude and direction are obtained as:

$$G = \sqrt{G_x^2 + G_y^2}, \quad \theta = \arctan\left(\frac{G_y}{G_x}\right) \quad (5)$$

673 To suppress high-frequency noise, we apply a Gaussian blur with a (9×9) kernel before edge
 674 detection. The Canny operator is then applied with lower and upper thresholds set to 10 and
 675 20, respectively.

676 Following edge extraction, we adopt a contour-based shape matching method to assess
 677 similarity between detected endpoints in the target (or distractor) image and a pre-defined
 678 template. From each binary edge map, external contours are identified by locating the
 679 outermost connected boundary points. Each contour is simplified to preserve essential
 680 geometric properties while reducing data redundancy. Let \mathcal{C}_1 and \mathcal{C}_2 denote the contour sets
 681 from images I_1 and I_2 , respectively. To enhance robustness against noise, only the most
 682 prominent contour—defined as the one enclosing the largest area—is selected from each set:

$$C_1^* = \arg \max_{C \in \mathcal{C}_1} \text{Area}(C), \quad C_2^* = \arg \max_{C \in \mathcal{C}_2} \text{Area}(C) \quad (6)$$

683 Shape similarity is quantified using a moment-based descriptor derived from the seven Hu
 684 moments H_1, H_2, \dots, H_7 , which are invariant to translation, rotation, and scaling. The similarity
 685 score between the selected contours is defined as

$$D(C_1^*, C_2^*) = \sum_{i=1}^7 \left| \frac{1}{\log H_i^{(1)}} - \frac{1}{\log H_i^{(2)}} \right| \quad (7)$$

686 where $H_i^{(1)}$ and $H_i^{(2)}$ are the i -th Hu moments computed from C_1^* and C_2^* , respectively. Lower
 687 value of D indicates greater shape similarity.

688 Comparative evaluations with alternative edge-detection and contour-matching strategies
 689 (Note S10, Supporting Information) confirmed that the proposed combination yields the most
 690 reliable and computationally lightweight performance, meeting the operational constraints of
 691 real-time WCE navigation. [To ensure practical reliability in diverse clinical settings, we further
 692 evaluated the module's performance under non-ideal conditions, including moderate off-axis
 693 tilts and dim-light environments. The module demonstrates consistent landmark identification
 694 and depth calculation accuracy even when features are positioned far from the optical center
 695 \(Note S12, Supporting Information\).](#)

696 **3D Landmark Localization**

697 To determine the precise three-dimensional coordinates of a navigation endpoint, we employed
 698 a deep neural network-based monocular depth estimation model, DispNet, configured within

699 the Endo-SfMLearner framework⁵⁰. The model receives a single monocular RGB frame I_i
700 captured by the capsule camera and outputs a disparity map D_i , effectively augmenting the
701 original image with per-pixel depth information. The encoder-decoder architecture of DispNet
702 (Note S11, Supporting Information) was selected for its ability to produce dense and smooth
703 disparity maps under the low-texture, specular, and variable-illumination conditions typical of
704 endoscopic imaging. The network was trained using a loss function that integrates geometry
705 consistency, brightness-aware photometric matching, and smoothness constraints, jointly
706 optimised to improve both the accuracy and robustness of estimated disparities.

707 Geometry Consistency Loss (L_{GC}) enforces agreement between forward-warped disparities
708 D_i^w and corresponding backward disparities D_{i+1} , penalizing mismatches via the relative
709 difference metric

$$D_{\text{diff}}(p) = \frac{|D_i^w(p) - D_{i+1}(p)|}{D_i^w(p) + D_{i+1}(p)} \quad (8)$$

710 where p indexes pixels and $D_i^w(p)$ denotes the disparity at p in frame i forward warped into
711 frame $i + 1$. The loss is averaged over all pixels in the set P :

$$L_{GC} = \frac{1}{|P|} \sum_{p \in P} D_{\text{diff}}(p) \quad (9)$$

712 Brightness-Aware Photometric Loss (L_p) measures the residual between the reference image
713 I_i and a synthesized view (\hat{I}_i) after applying an affine brightness correction $T_b(\cdot)$, thereby
714 compensating for lighting variation. This term blends an ℓ_1 photometric error with a structural
715 similarity index (SSIM) term for enhanced perceptual robustness:

$$L_p = \frac{1}{|P|} \sum_{p \in P} \left(\|T_b(\hat{I}_i)(p) - I_i(p)\|_1 + \lambda_1 \frac{1 - \text{SSIM}(T_b(\hat{I}_i), I_i)}{2} \right) \quad (10)$$

716 where

$$T_b(\hat{I}_i) = a_i \hat{I}_i + c_i w, \quad M = 1 - D_{\text{diff}} \quad (11)$$

717 and λ_1 controls the SSIM contribution.

718 Smoothness Loss (L_s) regularises disparity values to promote spatial coherence while
719 preserving depth discontinuities along image edges. By weighting the disparity gradient $\nabla D(p)$
720 by an exponential function of the image gradient magnitude, the loss emphasises smoothness
721 in low-texture regions:

$$L_s = \sum_{p \in P} (e^{-|\nabla I(p)|} \cdot |\nabla D(p)|)^2 \quad (12)$$

722 The total loss is a weighted sum of these three terms:

$$L = \alpha_p L_p + \beta_s L_s + \gamma_{GC} L_{GC} \quad (13)$$

723 where α_p , β_s , and γ_{GC} balance the contributions of photometric, smoothness, and geometric
724 consistency constraints.

725 This composite formulation ensures that depth estimates are accurate, geometrically consistent,
726 and resilient to lighting variability. In our implementation, we employed a pre-trained DispNet

727 model, achieving a root mean square error (RMSE) of 0.37 cm on a desktop platform equipped
728 with a 32 GB AMD Ryzen 5 5600 six-core processor and an NVIDIA GeForce RTX 4060 GPU,
729 meeting real-time inference requirements for intraoperative navigation.

730 **EndoSLAM Dataset**

731 DispNet was pre-trained using data from the EndoSLAM dataset ⁵⁰, a large-scale monocular
732 depth estimation benchmark comprising 42,700 frames. The dataset integrates ex-vivo porcine
733 gastrointestinal organ recordings with synthetic capsule endoscopy sequences generated in the
734 Unity simulation environment, covering anatomical regions including the colon, stomach, and
735 small intestine. Each sequence includes pose-aligned depth maps, obtained via a combination
736 of computed tomography (CT) imaging and high-accuracy 3D scanning, alongside
737 corresponding video streams captured by both capsule endoscopy and conventional endoscopic
738 cameras.

739 Although labelled and calibrated, EndoSLAM is designed primarily for unsupervised
740 monocular depth estimation, where models learn depth indirectly from photometric
741 consistency between consecutive frames rather than from direct pixel-wise supervision. This
742 structure facilitates domain adaptation, as the dataset supports the joint evaluation of 3D
743 reconstruction and simultaneous localisation and mapping (SLAM) techniques in both
744 synthetic and physical environments. Its composition is particularly suited to simulation-to-real
745 transfer learning, enabling pre-trained models to transfer effectively to real-world endoscopic
746 imagery for robust geometry prediction.

747 **Training Approach and workflow**

748 For AL-DRL training, we employed the PPO algorithm implemented within the Unity
749 ML-Agents framework. ML-Agents integrates natively with the Unity 3D simulation engine,
750 supporting both headless (terminal-based) training and interactive visual rendering, thereby
751 enabling flexible control over computational resources and debugging. The hyperparameters
752 were tuned (Table 1) to improve capsule performance within a given stomach model. For
753 landmark-guided training, the two anatomical endpoints were first localised using the Image
754 Processing Module. The agent was then trained using a reward function specifically formulated
755 for the offline AL-DRL framework, allowing precise evaluation of landmark-guided WCE
756 navigation performance in the anatomical model. In this setup, the observation vector—which
757 serves as the direct input to the DRL agent’s policy network—is composed of the 3D
758 coordinates of the target landmarks, combined with real-time spatial and temporal information
759 from WCE and EPM. This custom reward function, defined via C# script within the Unity
760 environment, then uses this real-time information to calculate the DRL reward for training.

761

```

762 Algorithm 1 Agent training in each episode
763 Require:
764 Position of the capsule and EPM, Velocity of capsule, Rotation of camera mounted on capsule,
765 Coordinates of two end points of stomach
766 Ensure:
767 Reward of every step
768 Initialization:
769     Capsule angular velocity  $\leftarrow 0$ , Capsule velocity  $\leftarrow 0$ ;
770     EPM angular velocity  $\leftarrow 0$ , EPM velocity  $\leftarrow 0$ ;
771     EPM and capsule position  $\leftarrow$  Initial position;
772     End1, End2  $\leftarrow$  Edge-contour-depth fusion module, Target end  $\leftarrow$  End2;
773     Boundary reward  $\leftarrow 0$ ;
774     End reward  $\leftarrow 0$ , Total reward[]  $\leftarrow 0$ ;
775     Coverage reward[]  $\leftarrow 0$ , Track reward[]  $\leftarrow 0$ ;
776 while step  $\leq 10000$  do
777     Collect Observations:
778         Capsule velocity and rotation  $\leftarrow$  Unity;
779         EPM position and rotation relative to the capsule  $\leftarrow$  Unity;
780         End points position relative to the capsule  $\leftarrow$  Unity;
781     Get action from PPO:
782         EPM velocity  $\leftarrow$  Neural network;
783         EPM angular velocity  $\leftarrow$  Neural network;
784         EPM position  $\leftarrow$  EPM position + EPM velocity  $\times \Delta t$ ;
785         EPM rotation  $\leftarrow$  EPM rotation + EPM angular velocity  $\times \Delta t$ ;
786         Capsule position and rotation  $\leftarrow$  Dipole model;
787     Calculate reward:
788         Coverage diff  $\leftarrow$  |Current coverage – Previous coverage|;
789     if Coverage diff  $> 0.03$  then
790         Coverage reward[step]  $\leftarrow$  Coverage diff  $\times 20$ ;
791     else
792         Coverage reward[step]  $\leftarrow -1$ ;
793     end if
794     if |Capsule position – End2|  $\leq 0.3$  and Target end == End2 then
795         End reward  $\leftarrow 3$ ;
796         Target end  $\leftarrow$  End1;
797     else if |Capsule position – End1|  $\leq 0.3$  and Target end == End1 then
798         End reward  $\leftarrow 3$ ;
799         Target end  $\leftarrow$  End2;
800     end if

```

```
801         Track reward[step] ← -|Capsule position – Target end| + End reward;
802     if |Capsule position – EPM position| > 4 then
803         Boundary reward ← -50;
804         End episode
805     end if
806
806     Total reward[step] ← Coverage reward[step] + Track reward[step] + Boundary
807     reward.
808 end while


---


809
```

810 **Table 1**
 811 *Hyperparameters in the PPO, SAC, and DQN Algorithms*

Parameter	PPO	SAC	DQN
Trainer	PPO	SAC	DQN
Batch Size	1024	1024	64
Buffer Size	4096	512000	50000
Hidden Units	128	128	128
Learning Rate	0.001	0.0005	0.0003
Learning Rate Schedule	Linear	Linear	constant
Max Steps	3000000	3000000	3000000
Memory Size	512	512	10000
Num Layers	2	2	2
Time Horizon	1024	1024	1024
Sequence Length	64	64	N/A
Summary Freq	10000	10000	1000
Gamma	0.99	0.99	0.99
Num Epoch	5	N/A	N/A
Lambda	0.95	N/A	N/A
Epsilon	0.2	N/A	1.0
Beta	0.005	N/A	N/A
Tau	N/A	0.005	0.005
Exploration Schedule	N/A	N/A	Linear
Exploration Initial Eps	N/A	N/A	0.8
Exploration Final Eps	N/A	N/A	0.05

812

813

814 **References**

- 815 1 Swain, P. Wireless capsule endoscopy. *Gut* **52**, iv48-iv50 (2003).
- 816 2 Fireman, Z. *et al.* Diagnosing small bowel Crohn's disease with wireless capsule
817 endoscopy. *Gut* **52**, 390-392 (2003).
- 818 3 Schwartz, G. D. & Barkin, J. S. Small-bowel tumors detected by wireless capsule
819 endoscopy. *Digestive diseases and sciences* **52**, 1026-1030 (2007).
- 820 4 Cao, Q. *et al.* Robotic wireless capsule endoscopy: recent advances and upcoming
821 technologies. *Nature Communications* **15**, 4597 (2024). [https://doi.org/10.1038/s41467-](https://doi.org/10.1038/s41467-024-49019-0)
822 [024-49019-0](https://doi.org/10.1038/s41467-024-49019-0)
- 823 5 Orłowska, J., Jarosz, D., Pachlewski, J. & Butruk, E. Malignant transformation of benign
824 epithelial gastric polyps. *American Journal of Gastroenterology (Springer Nature)* **90**
825 (1995).
- 826 6 Jung, J. T. Gastric polyps and protruding type gastric cancer. *Clinical Endoscopy* **46**, 243
827 (2013).
- 828 7 Imura, J. *et al.* Malignant transformation of hyperplastic gastric polyps: An
829 immunohistochemical and pathological study of the changes of neoplastic phenotype.
830 *Oncology letters* **7**, 1459-1463 (2014).
- 831 8 Shaib, Y. H., Ruge, M., Graham, D. Y. & Genta, R. M. Management of gastric polyps:
832 an endoscopy-based approach. *Clinical gastroenterology and hepatology* **11**, 1374-1384
833 (2013).
- 834 9 Morais, D. J., Yamanaka, A., Zeitune, J. M. R. & Andreollo, N. A. Gastric polyps: a
835 retrospective analysis of 26,000 digestive endoscopies. *Arquivos de gastroenterologia* **44**,
836 14-17 (2007).
- 837 10 Moglia, A., Menciacchi, A., Dario, P. & Cuschieri, A. Capsule endoscopy: progress update
838 and challenges ahead. *Nature Reviews Gastroenterology & Hepatology* **6**, 353-361 (2009).
- 839 11 Geropoulos, G., Aquilina, J., Kakos, C., Anestiadou, E. & Giannis, D. Magnetically
840 controlled capsule endoscopy versus conventional gastroscopy: a systematic review and
841 meta-analysis. *Journal of Clinical Gastroenterology* **55**, 577-585 (2021).
- 842 12 Jiang, X., Pan, J., Li, Z.-S. & Liao, Z. Standardized examination procedure of
843 magnetically controlled capsule endoscopy. *VideoGIE* **4**, 239-243 (2019).
- 844 13 Shamsudhin, N. *et al.* Magnetically guided capsule endoscopy. *Medical physics* **44**, e91-
845 e111 (2017).
- 846 14 Slawinski, P. R., Obstein, K. L. & Valdastrì, P. Capsule endoscopy of the future: What's
847 on the horizon? *World journal of gastroenterology: WJG* **21**, 10528 (2015).
- 848 15 Carpi, F. & Pappone, C. Magnetic maneuvering of endoscopic capsules by means of a
849 robotic navigation system. *IEEE transactions on biomedical engineering* **56**, 1482-1490
850 (2009).

- 851 16 Ciuti, G. *et al.* Frontiers of robotic endoscopic capsules: a review. *Journal of micro-bio*
852 *robotics* **11**, 1-18 (2016).
- 853 17 Sun, Y. *et al.* Magnetically driven capsules with multimodal response and
854 multifunctionality for biomedical applications. *Nature Communications* **15**, 1839 (2024).
- 855 18 Hoang, M. C. *et al.* Independent electromagnetic field control for practical approach to
856 actively locomotive wireless capsule endoscope. *IEEE Transactions on Systems, Man,*
857 *and Cybernetics: Systems* **51**, 3040-3052 (2019).
- 858 19 Mahoney, A. W. & Abbott, J. J. Five-degree-of-freedom manipulation of an untethered
859 magnetic device in fluid using a single permanent magnet with application in stomach
860 capsule endoscopy. *The International Journal of Robotics Research* **35**, 129-147 (2016).
- 861 20 McGoran, J. J. *et al.* Miniature gastrointestinal endoscopy: Now and the future. *World*
862 *journal of gastroenterology* **25**, 4051 (2019).
- 863 21 Yuan, S. *et al.* Motor-free telerobotic endomicroscopy for steerable and programmable
864 imaging in complex curved and localized areas. *Nature Communications* **15**, 7680 (2024).
- 865 22 Mao, L. *et al.* Magnetic steering continuum robot for transluminal procedures with
866 programmable shape and functionalities. *Nature communications* **15**, 3759 (2024).
- 867 23 Wang, H., Cui, J., Tian, K. & Han, Y. Three-degrees-of-freedom orientation manipulation
868 of small untethered robots with a single anisotropic soft magnet. *Nature Communications*
869 **14**, 7491 (2023).
- 870 24 Mateen, H., Basar, R., Ahmed, A. U. & Ahmad, M. Y. Localization of wireless capsule
871 endoscope: A systematic review. *IEEE Sensors Journal* **17**, 1197-1206 (2017).
- 872 25 Ciuti, G., Menciassi, A. & Dario, P. Capsule endoscopy: from current achievements to
873 open challenges. *IEEE reviews in biomedical engineering* **4**, 59-72 (2011).
- 874 26 Zhang, H. *et al.* Toward automatic stomach screening using a wireless magnetically
875 actuated capsule endoscope. *IEEE Transactions on Medical Robotics and Bionics* **6**, 512-
876 523 (2024).
- 877 27 Jain, S. *et al.* A deep CNN model for anomaly detection and localization in wireless
878 capsule endoscopy images. *Computers in Biology and Medicine* **137**, 104789 (2021).
- 879 28 Turan, M., Shabbir, J., Araujo, H., Konukoglu, E. & Sitti, M. A deep learning based fusion
880 of RGB camera information and magnetic localization information for endoscopic capsule
881 robots. *International journal of intelligent robotics and applications* **1**, 442-450 (2017).
- 882 29 Cui, G. *et al.* Soft objects grasping evaluation using a novel VCFN-YOLOv8 framework.
883 *Biomimetic Intelligence and Robotics* **5**, 100232 (2025).
884 <https://doi.org/https://doi.org/10.1016/j.birob.2025.100232>
- 885 30 Turan, M. *et al.* Learning to navigate endoscopic capsule robots. *IEEE Robotics and*
886 *Automation Letters* **4**, 3075-3082 (2019).

- 887 31 Hanscom, M. & Cave, D. R. Endoscopic capsule robot-based diagnosis, navigation and
888 localization in the gastrointestinal tract. *Frontiers in Robotics and AI* **9**, 896028 (2022).
- 889 32 Zhang, Y., Bai, L., Liu, L., Ren, H. & Meng, M. Q. H. in *2022 IEEE International*
890 *Conference on Robotics and Biomimetics (ROBIO)*. 01-06 (IEEE).
- 891 33 İncetan, K. *et al.* VR-Caps: a virtual environment for capsule endoscopy. *Medical image*
892 *analysis* **70**, 101990 (2021).
- 893 34 Di Natale, M. R. *et al.* Functional and anatomical gastric regions and their relations to
894 motility control. *Neurogastroenterology & Motility* **35**, e14560 (2023).
895 <https://doi.org/https://doi.org/10.1111/nmo.14560>
- 896 35 Upadhyay, A. K. & Bhandari, A. K. Advances in deep learning models for resolving
897 medical image segmentation data scarcity problem: A topical review. *Archives of*
898 *Computational Methods in Engineering* **31**, 1701-1719 (2024).
- 899 36 Bansal, M. A., Sharma, D. R. & Kathuria, D. M. A systematic review on data scarcity
900 problem in deep learning: solution and applications. *ACM Computing Surveys (Csur)* **54**,
901 1-29 (2022).
- 902 37 Su, S. *et al.* AMagPoseNet: Real-Time Six-DoF Magnet Pose Estimation by Dual-Domain
903 Few-Shot Learning From Prior Model. *IEEE Transactions on Industrial Informatics* **19**,
904 9722-9732 (2023). <https://doi.org/10.1109/TII.2022.3233675>
- 905 38 Wang, Y., He, H. & Tan, X. in *Uncertainty in artificial intelligence*. 113-122 (PMLR).
- 906 39 Fan, J., Wang, Z., Xie, Y. & Yang, Z. in *Learning for dynamics and control*. 486-489
907 (PMLR).
- 908 40 Haarnoja, T., Zhou, A., Abbeel, P. & Levine, S. in *International conference on machine*
909 *learning*. 1861-1870 (Pmlr).
- 910 41 Zhang, J., Zhang, Z., Han, S. & Lü, S. Proximal policy optimization via enhanced
911 exploration efficiency. *Information Sciences* **609**, 750-765 (2022).
- 912 42 Huang, S., Xiao, Z., Zheng, M. & Shi, W. Hierarchical reinforcement learning for
913 enhancing stability and adaptability of hexapod robots in complex terrains. *Biomimetic*
914 *Intelligence and Robotics* **5**, 100231 (2025).
915 <https://doi.org/https://doi.org/10.1016/j.birob.2025.100231>
- 916 43 Rizk, M. K. *et al.* Quality indicators common to all GI endoscopic procedures. *Official*
917 *journal of the American College of Gastroenterology| ACG* **110**, 48-59 (2015).
- 918 44 Niv, Y. & Niv, G. Capsule endoscopy: role of bowel preparation in successful
919 visualization. *Scandinavian journal of gastroenterology* **39**, 1005-1009 (2004).
- 920 45 Rey, J.-F. *et al.* Feasibility of stomach exploration with a guided capsule endoscope.
921 *Endoscopy* **42**, 541-545 (2010).

- 922 46 Yao, T. *et al.* Sim2real learning with domain randomization for autonomous guidewire
923 navigation in robotic-assisted endovascular procedures. *IEEE Transactions on*
924 *Automation Science and Engineering* (2025).
- 925 47 Zhao, B. & Li, Y. Model-free adaptive dynamic programming based near-optimal
926 decentralized tracking control of reconfigurable manipulators. *International Journal of*
927 *Control, Automation and Systems* **16**, 478-490 (2018).
- 928 48 Su, S., Lan, H., Gao, H., Dai, H. & Ren, H. Enhancing Anti-Interference of Magnetic
929 Tracking: A MagRobustNet-Based Framework With Self-Supervised Anomaly Detection
930 and Measurements Recovery. *IEEE Transactions on Industrial Informatics* **21**, 781-790
931 (2025). <https://doi.org/10.1109/TII.2024.3461774>
- 932 49 Liao, Z. *et al.* Accuracy of Magnetically Controlled Capsule Endoscopy, Compared With
933 Conventional Gastroscopy, in Detection of Gastric Diseases. *Clinical Gastroenterology*
934 *and Hepatology* **14**, 1266-1273.e1261 (2016).
935 <https://doi.org/https://doi.org/10.1016/j.cgh.2016.05.013>
- 936 50 Ozyoruk, K. B. *et al.* EndoSLAM dataset and an unsupervised monocular visual odometry
937 and depth estimation approach for endoscopic videos. *Medical image analysis* **71**, 102058
938 (2021).

939

940 **Animal studies**

941 In vivo experiment (Figure S18, Supporting Information) was conducted on Yorkshire swine
942 (n=1, weight 30 kg), obtained from the Shandong Quality Inspection Center for Medical
943 Devices. All experimental procedures and animal handling protocols were rigorously reviewed
944 and approved by the Qilu Hospital of Shandong University. Prior to the experiment, the animal
945 was fasted for 12 hours to ensure a clear gastric view. Anesthesia was induced via an
946 intramuscular injection of ketamine (25 mg/kg) and subsequently maintained through
947 continuous intravenous infusion of propofol (800 µg/kg/h) to ensure stable sedation throughout
948 the navigation procedure. Vital signs were continuously monitored to maintain physiological
949 stability.

950

951 **Data Availability Statement**

952 The main data supporting the results in this study are available within the paper and its
953 supporting information.

954

955 **Acknowledgements**

956 This work was supported in part by NSFC Young Scientists Fund - Category A T252500134,
957 Hong Kong Research Grants Council (RGC) Collaborative Research Fund (CRF C4026-21GF),
958 and General Research Fund (GRF 14216022, 14204524, 14203323, 14206125), NSFC/RGC

959 Joint Research Scheme N_CUHK420/22; Guangdong Basic and Applied Basic Research
960 Foundation (GBABF) \#2021B1515120035.

961

962 **Conflict of Interest**

963 The authors declare no conflict of interest.