Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

SYNEX: Perception-to-valve autonomous excavation with coverage planning and worker protection

Mehdi Heydari Shahna
Tampere University
Seihun Kim
Korea University
Soyi Jung
Ajou University
Soohyun Park
Sookmyung Women's University
Jouni Mattila
Tampere University
Joongheon Kim
joongheon@korea.ac.kr

Korea University

Research Article

Keywords:

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11823481/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

SYNEX: Perception-to-valve autonomous excavation with coverage planning and worker protection

Mehdi Heydari Shahna[a], Seihun Kim[b], Soyi Jung[c,*], Soohyun Park[d,*], Jouni Mattila[a] and Joongheon Kim[e,*]

[a]*Faculty of Engineering and Natural Sciences, Tampere University, Tampere, Finland*
[b]*Department of Defense Convergence Technology, Korea University, Seoul, Republic of Korea*
[c]*Department of Electrical and Computer Engineering, Ajou University, Suwon, Republic of Korea*
[d]*Division of Computer Science, Sookmyung Women's University, Seoul, Republic of Korea*
[e]*Department of Electrical and Computer Engineering, Korea University, Seoul, Republic of Korea*

ARTICLE INFO

Keywords:
Autonomous excavation
Construction robotics
Coverage planning
Control barrier functions
Speed and separation monitoring
Digital twin

ABSTRACT

SYNEX connects site perception to hydraulic valve commands for multi-pit footing excavation. We evaluated it on a digital twin of a 30 t excavator whose front and hydraulics were calibrated against field records. The system combines stable-wall layered coverage, synchronized planning, two control barrier function layers, and a worker separation monitor based on stopping bounds measured in simulation. In ten missions with perception, 99 of 100 pits met the grade tolerance at all 289 grid points, including 49 of 50 in held-out missions. Nine of the ten missions completed the work order. In 40 prescribed walking-worker scenarios, repeated with exact states, perception, and a perception blackout, no clearance below 1.5 m was observed while the excavator approached a worker faster than 0.05 m/s. These results describe performance in simulation under the tested conditions.

1. Introduction

Hydraulic excavators repeat a linked sequence of construction tasks: travel to a working position, dig to a specified grade, carry the material and unload it into a truck or onto a spoil area. Automating this sequence requires more than a successful bucket trajectory. The system must interpret the work order, locate the excavation and unloading areas, coordinate its motion with people and equipment, and establish that the required earthwork was actually completed. Reviews in *Automation in Construction* organize the problem around sensing, perception, decision-making, planning and action [1] and identify the integration of individual stages into a complete workflow as an unresolved challenge [2].

Three coupled difficulties stand in the way of such an integration. First, hydraulic commands are not realized immediately or exactly: valve commands act after a dead time of 0.1 s to 0.2 s, field records of the machine studied here show a valve deadzone of 33 % to 41 % of the duty range, and pressure dynamics and shared pump flow limit simultaneous motion [3]. A planner can therefore request a collision-free motion that the actuators cannot follow. Second, a pit is not finished when most of its soil is removed: the work order asks for every point of the design surface within a grade tolerance, which depends on how the soil behaves at the pit walls, on how accurately the bucket follows its passes, and on whether the excavated material can be placed. Third, people walk on the site; a 30 t machine needs seconds and metres to stop, and the people may be hidden from the machine's sensors by the truck, the boom or the machine itself. Because such a system is developed and evaluated in simulation, the evidence must in addition show which mechanisms of the simulation drive its outcomes and how far the simulation is checked against the real machine.

We present *SYNEX*, an autonomous excavation system that connects site perception to the valve commands of the machine and is designed to complete multi-pit work orders to an every-point grade criterion. It is evaluated on a digital twin of a 30 t Doosan DX300LC-5, whose front and its hydraulics are calibrated against field records, in a construction-site simulation with ten footing pits, a dump truck, a spoil area, static obstacles and walking workers. The contributions are:

*Corresponding author
sjung@ajou.ac.kr (S. Jung); soohyun.park@sookmyung.ac.kr (S. Park); joongheon@korea.ac.kr (J. Kim)

1. *A perception-to-valve architecture with synchronized two-layer safety filtering.* LiDAR mapping and camera segmentation provide the terrain, obstacles and object tracks, confirm the marked pits and locate the marked spoil area; a synchronized planner and two control-barrier-function layers, one on the motion reference and one on the valve duties with dead-time prediction and a braking credit, reduce the violations of the joint, self-collision and ground boundaries in a randomized benchmark to a small residual (0.06 % of the control steps, all of them joint and ground steps), and the physics audit of the full-twin experiments checks the safety-critical boundaries (self-collision, obstacles, workers, footprint) with the true geometry.
2. *A construction-task layer designed to complete pits to grade.* Stable-wall layered coverage keeps the pit walls below the failure step of the soil model and every pass within the bucket's capacity; pass margins follow from the cutting process; near the far wall the drag keeps the bucket open, so that its heel does not hold the teeth above the grade; every grade strip ends with a thin pass that clears the near rows hidden from the LiDAR; above the grade layer, a small load kept in the bucket between passes is released only over the middle of the next pass, and on the grade layer every load is carried out of the pit; every pass is aligned with its strip line before the teeth enter the soil, with a swing actuation fine enough for that alignment; dump points are planned within the reach of the curled bucket and spread over the spoil area; full trucks are exchanged. Five liveness rules release the machine from states in which the safety layers, the tracker and the task layer would otherwise hold it indefinitely. With perception in the loop, 99 of 100 pits of ten ten-pit missions reached the grade at every point, 49 of the 50 pits of the five missions on held-out scenario seeds among them, and nine of the ten missions completed the whole work order; the pit that missed the grade is traced to its mechanism.
3. *Worker protection with identified stopping bounds and an explicit operating envelope.* A separation monitor on the swept stopping region of the machine protects tracked workers, persons detected by near-field scanners on the undercarriage and workers lost from view, through their reachable sets, by stopping the machine or limiting the speed of its planned motion. Its stopping bounds are evaluated at an effective speed that accounts for the motion reference leading the hydraulics; identified and checked on separate sets of controlled stops, they held for every stop of a third set measured after they were fixed. The design rule is tested inside and outside the stated envelope, with machine-induced and worker-induced clearance losses and three-dimensional contact separated: in 40 prescribed encounters inside the envelope, no clearance below the standoff was observed during machine approach faster than 0.05 m/s.
4. *An evaluation that exposes the mechanisms behind the outcomes.* Design ablations remove one element of the task layer at a time on single pits, a physics audit checks the safety-critical boundaries with the true geometry in every run, the twin is compared with the manufacturer's data on working ranges, digging forces, swing and travel speeds and bucket capacity, and the sensitivity of pit completion to the soil parameters and to the pump capability is measured; the claims are scoped accordingly.

Learned residual policies on the navigation and excavation priors were trained with a performance-based checkpoint gate and evaluated on held-out cohorts; they gave no practical advantage over the priors, which the evaluated system therefore uses (Section 10.7). Code, models, configurations, result files and a supplementary video accompany the manuscript.

2. Related Work

Autonomous excavation systems and task planning. Eraliev et al. [1] review the sensing, perception, decision, planning and action stages of autonomous excavators, and Feng et al. [2] emphasize the remaining gap between subsystem advances and coordinated operation in repetitive earthmoving. Complete systems have demonstrated integrated perception, planning and control on real machines, including the HEAP walking excavator, which also digs trenches autonomously [4], and an excavator for material loading [5]. Task planners divide the work area into digging units and schedule the excavator's paths from a site model [6], and planning for autonomous excavation has been combined with force-based bucket control to reach a desired terrain shape [7]. Our task layer addresses a narrower question that these systems leave open: which planning and execution rules let every point of a pit reach its grade tolerance when the soil fails at steep

walls, the passes must follow their strips to the boundary of the pit and the dump place must stay within reach. Our study is in simulation; it does not provide comparable field evidence.

Learning for excavation. Hydraulic nonlinearities, pressure dynamics and flow limits make excavator control difficult [3]. Reinforcement learning has produced arm controllers with data-driven actuator models [8] and soil-adaptive excavation trajectories [9]; Cho et al. [10] propose a collaborative learning architecture for excavator planning and execution, and Molaei et al. [11] capture rocks with a standard excavator through model-free reinforcement learning. We use learning only as residual corrections on structured priors and test, under a performance-based checkpoint gate, whether it improves on them.

Safety filters and worker protection. A control barrier function (CBF) expresses a safety requirement as an inequality such as a nonnegative clearance, and a quadratic program can modify a nominal command to satisfy it with minimal change [12, 13]; exponential and high-order CBFs address constraints of higher relative degree [14, 15], and input-delay methods predict the state over the delay [16]. Rate-consistent filtering accounts for limits on reference tracking [17]; we use that reference filter and add an actuator filter for hydraulic dead time, lag and deadzone. Shields and CBF layers have been combined with learning to constrain exploration [18, 19]. For people near machines, speed and separation monitoring keeps a protective separation distance that covers the person's approach during the machine's reaction and stopping time [20], with the walking approach speed of ISO 13855 [21]; for earth-moving machinery, ISO 21815-1 specifies collision warning and avoidance by speed reduction or motion inhibit [22]. Proximity warning systems alert workers and operators when a worker comes within a set distance of the equipment [23], and radio-frequency tags worn by the workers can inhibit the swing and travel of an excavator through its pilot circuit [24]. Lin et al. [25] monitor the safety of worker–robot collaboration in construction with a digital twin, and Fu et al. [26] generate excavation trajectories for uncrewed mining excavators with a digital twin. Our monitor follows the structure of the protective separation distance, but bounds the machine's stopping travel on its swept stopping region and identifies its stopping bounds on the twin with a held-out check.

Simulation, mapping and soil. We compile the closed-chain front dynamics with RoboCompiler [27, 28] for a batched CPU twin; parallel simulation [29, 30] supports policy optimization [31, 32] with privileged critics [33–35], and residual learning [36, 37] restricts the initial search to the prior's neighbourhood. Mapping with truncated and Euclidean signed distance fields (TSDF/ESDF) [38, 39], range-image clustering [40] and encoder–decoder segmentation [41] supply the perceived world. Earthmoving force models [42, 43], multiscale soil parameters [44] and heightfield soil representations [45] supply the soil–tool interaction.

3. Machine, Twins and Construction Site

3.1. Machine and full twin

The machine is a Doosan DX300LC-5 (30.9 t; boom 6.245 m, arm 3.1 m, bucket 1.27 m^3). Its simulation model, the *full twin*, is a model of the machine, not a live counterpart of it. Its front (boom, arm, bucket and the bucket four-bar linkage) is modelled in MuJoCo [46] as a closed chain; the electro-hydraulics comprise the pilot stage (dead time, pressure-reducing valve maps, rise and fall lags), main-spool lags, meter-in and meter-out orifices, two pressure-controlled pumps with boom/arm confluence and boom priority, arm regeneration, chamber compressibility, reliefs, end-of-stroke cushions and seal friction. The input is the valve duty $d \in [-100, 100]^6$ of the six functions (boom, arm, bucket, swing, left and right travel), as on the machine's controller. The front and its hydraulics were calibrated against two field sessions recorded on the machine (single-joint duty ramps and multi-joint motions in free air). In open-loop replay of the recorded duties over windows of 7 s to 13 s, each started from the measured state (eight windows per session), the twin reproduces the joint angles with RMS errors of 0.16, 0.41 and 0.33 ° (boom, arm, bucket; session 1) and 0.44, 0.34 and 0.36 ° (session 2) and the chamber pressures within 6 bar to 13 bar RMS; replayed over a whole session (72 s) without re-initialization, the errors are 0.20, 1.03 and 0.75 ° and 0.48, 0.61 and 0.57 °. A parameter set fitted only to session 1 predicts session 2 with 0.77, 2.49 and 2.40 ° (windows). On the machine, motion starts at 33 % to 41 % duty (valve deadzone); the twin reproduces these thresholds within 1.9 duty points. The calibrated parameter set (pilot-stage dead times, pressure maps and lags, and the hydraulic and friction parameters of the front) and the replay summaries are part of the code archive. The swing and travel drives were not exercised in the field records; they are sized from the manufacturer's specification, with assumed valve curves; the soil model is not calibrated. The swing circuit has an anti-reaction function

that connects the motor ports for 1 s after the swing pilots close, which damps the rebound of a stopping swing; afterwards the closed valve holds the upper structure against side loads. Section 10.1 compares the twin with the manufacturer's specification sheet (working ranges, digging forces, swing and travel speeds), and Section 10.4 examines the soil model.

The soil is a heightfield (0.1 m cells) with bank and loose layers. The cutting resistance follows the fundamental earthmoving equation [42, 43] with the parameters of a dirt from [44] (bank density $1{,}600\,\mathrm{kg/m^3}$, friction angle 37°, cohesion 4.9 kPa, swell factor 1.25); the bucket shell rests on a bearing-capped Winkler support; soil that the shell overlaps without taking it in is moved to the first cells outside the bucket's footprint along the shell's motion and relaxed to its angle of repose [45] (35° loose), spilled soil flies ballistically, and the soil mass is conserved to round-off. Cut walls of bank soil stand up to a height step of $\tan(70°) \times 0.1\,\mathrm{m}$ between neighbouring cells and fail beyond it. This wall rule is a cellular rule of the soil model, independent of the cutting parameters: for the cohesion and friction angle above, the Mohr–Coulomb critical height of a vertical cut, $4c\tan(45° + \varphi/2)/\gamma$ with the unit weight γ of the bank soil, would be about 2.5 m. The twin therefore represents a soil whose cut walls do not stand vertically, as dry granular or fissured soil, and Section 10.4 varies the steepest standing wall. No digging data were recorded on the machine; the soil parameters are literature values, and Section 10.4 varies them.

3.2. Reduced-order batched twin

For training and for the control-layer ablations we compile the calibrated front with RoboCompiler [27, 28], which generates the spanning tree, the loop cuts and the reduced interfaces from the physical joint records. The reduced mass matrix of (swing, boom, arm, bucket), the gravity basis and the cylinder port maps are tabulated on a 5° grid of the front angles; they agree with MuJoCo's projection of the same closed chain to relative errors below 5×10^{-9} (gravity) and 4.8×10^{-4} (mass matrix). A C kernel integrates the reduced dynamics of a batch of machines with the hydraulics of the full twin (5 ms substeps). Driven open loop by the same 24 random duty sequences of 10 s, its joint angles deviate from the full twin's by $(0.21, 0.15, 0.40)°$ RMS; on one CPU core it runs 57,783 environment steps per second with 1024 environments, versus 118 control steps per second for the full twin. The batched twin uses exact states (no perception) and the travel model of the full twin; obstacles are analytic boxes and cylinders, workers are disks that walk on random paths, and the soil of each excavation zone is a height profile along the zone's radial axis with the constitutive content of the full twin's soil model (cutting resistance, scooping, retained capacity and spill).

3.3. Construction site

The site (Fig. 1) is a $44\,\mathrm{m}\times44\,\mathrm{m}$ soil heightfield within a coarser surrounding terrain surface; the simulated LiDARs return the soil surface wherever the soil heightfield exists. Ten footing pits P1–P10 of $1.6\,\mathrm{m}\times1.6\,\mathrm{m}$ with a design grade 0.6 m below the ground ($1.54\,\mathrm{m^3}$ each, $15.4\,\mathrm{m^3}$ in total) are marked with orange paint and two cones 2 m behind each pit, on level prepared ground along a building line 6 m beside the travel line; the machine excavates a pit from its *stance*, a position on the travel line abreast of the pit centre, so that the radial dig direction is parallel to one axis of the pit, with its swing axis about 6 m from the pit centre. The spoil of P1–P5 is loaded into a dump truck ($9.5\,\mathrm{m^3}$ of loose soil) that drives along a lane 6 m on the other side of the travel line; the spoil of P6–P10 goes to a blue-marked spoil area on the truck side, 5.5 m to 8.5 m from the travel line along P6–P10. Excavating a pit therefore involves swings of about 180°.

Static obstacles (light tower, pipe stack, site office, portable toilet, roller, pallets, rebar, jersey barriers, fence, cones) and three walking workers make the site cluttered and dynamic: W1 walks along the fence behind the stored material, the route of W2 crosses the travel corridor 12 m ahead of the machine's base and moves with the machine (from P6 on, where the walkable site ends less than the distance of the site rule beyond that line, it crosses 12 m behind the base), so that W2 keeps meeting the zone of the site rule below, and W3 walks along the barriers behind the truck lane. The site workers keep the site rule of an autonomous machine's work zone: they stay at least 12.5 m (surface to surface) from the machine's current outline. A worker whose next waypoint lies across this zone walks around it, or turns back where no detour exists, and steps back when the machine itself comes closer; where the walkable site ends, at the fence or the barriers, the worker steps back along that edge. A worker that cannot step back any farther, at the east end of the site, leaves the zone: it walks to the fence or the barrier line and along it past the machine, and resumes its route there (the exit rule; Section 9). The distance exceeds the separation that the monitor of Section 6.2

(a) Ten-location footing-excavation site

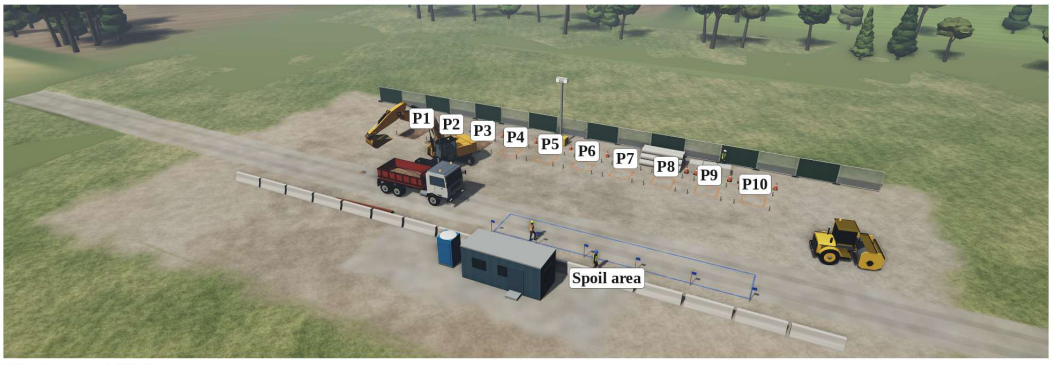

Mission time 1600.0 s

(b) Excavator dumping into the truck



Mission time 87.0 s

(c) Worker walking past the machine



Mission time 120.0 s

Figure 1: Presentation rendering of the construction site with recorded states of the ten-pit mission of scenario seed 401 (mission times under the panels). (a) The ten marked footing pits (P1–P10; orange outlines), the truck lane, the blue-marked spoil area for P6–P10 and the site obstacles. (b) The excavator dumping into the truck. (c) A site worker walking past the machine, outside the zone of the site rule. Machine poses and worker and truck positions are recorded; the walking gait is synthesized from the recorded positions and the soil surface is reconstructed for display; the terrain and scenery beyond the site are decorative. The presentation scene is separate from the scene of the sensor simulation used in the experiments.

needs ahead of a collision circle (one of the circles that represent the front, the upper structure and the tracks in the safety layers, Section 4) that moves toward a person at 3.5 m/s, the largest circle speed of the measured stops, rounded up (12.2 m), but not the separation at the largest effective speed of these stops (13.7 m at 4.1 m/s, while the motion reference led the hydraulics at the onset of a swing), so a worker who keeps the rule can still restrict a starting swing; closer to a person, the monitor limits or stops the motion. The site workers do not anticipate the machine's motion. The workers of the worker-protection experiments ignore this rule (Section 9). Two LiDARs (1° azimuth resolution, 60 m range, 2 cm noise, 1 % dropouts)—a 48-channel unit on the cab roof (−45° to 15°) that observes the dig area and a 32-channel unit on a mast on the right side of the deck (−25° to 15°)—, the four near-field scanners of Section 5.3 and three cab cameras (front, left, right; 240 × 180 pixels) form the sensor set.

4. Synchronized Planning and Two-Layer Safety Filters

Figure 2 shows the control stack. A *skill* (travel to a stance, or the dig–dump cycle of Section 7) issues normalized task-space intents $v \in [-1, 1]^6$ (rates of the swing angle ψ, the radial and vertical bucket-lip coordinates s, z and the bucket pitch ϕ, or the travel speed and yaw rate) at 20 Hz. The separation monitor of Section 6 may limit the speed of the planned motion or replace the intents by a controlled stop. The synchronized planner converts the intents into a smooth nominal coordinate rate u_n; the reference layer produces a filtered reference (u_r, a_r); the actuator controller proposes valve duties d_π; the actuator layer

filters them into the applied duties d. Both layers evaluate the same barrier functions. An *intent* specifies what a skill would like to do, a *reference* is the motion passed to the actuator controller and a *duty* is the final valve command; Table 1 lists the signals.

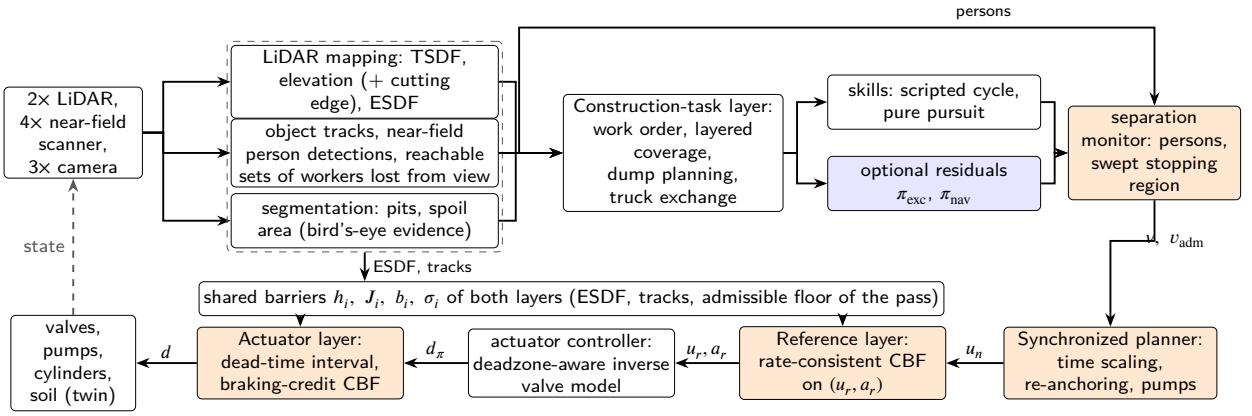

Figure 2: Perception-to-valve architecture. Safety and synchronization components are shaded orange, optional learned residuals blue; the evaluated system uses the scripted cycle, pure pursuit and the inverse valve model. The two safety layers share the barrier functions h_i, their gradients J_i, drift terms b_i and standoffs σ_i, computed from the perceived world (ESDF, object tracks, excavation zones); the separation monitor uses the tracked workers, the near-field person detections and the reachable sets of workers lost from view; the planner is time-scaled by the plan error, the pump capability and the admissible speed of the monitor.

Table 1
Signals of the control stack. A requested rate and a realized rate have the same units but differ because of hydraulic delay, loading and filtering.

Symbol	Meaning
ν	Dimensionless skill intent between -1 and 1.
u_n	Nominal coordinate rate of the planner.
(u_r, a_r)	Filtered reference rate and its acceleration, passed to the actuator controller.
u	Measured coordinate rate: four angular rates, travel speed and yaw rate.
d_π, d	Proposed and applied valve duties, in percent of the signed command range.
$w = V(d)$	Steady coordinate rate predicted by the identified valve map.
h_i, $\tilde{h}_i$	Modelled boundary value and requested-clearance margin $h_i - \sigma_i$.

4.1. Barrier functions

The coordinate-rate vector is $u = (\dot{\psi}, \dot{q}_1, \dot{q}_2, \dot{q}_3, v, \omega)$ (swing, boom, arm, bucket, travel speed, yaw rate). The modelled collision-free domain is described by $R = 18$ barrier rows $h_i(x,t) \geq 0$ of the machine state x (joint angles and base pose) and time, and the desired operating set by $\tilde{h}_i = h_i - \sigma_i \geq 0$ with standoffs σ_i: six joint limits (1° inside the cylinder end stops); two self-collision rows (spheres on the front against boxes of the cab, deck and boom foot, and against the track frames); two ground rows (the cutting-edge points and shell spheres of the bucket above the admissible floor, and the arm and boom spheres above the admissible floor); four static-obstacle rows (the most critical sphere–obstacle pairs, with distances from the ESDF of Section 5); three worker rows (the most critical sphere or footprint circle per worker); and one row for the base footprint during travel. The collision geometry consists of spheres on the front (on the bucket three cutting-edge points of radius 0.02 m, three lip spheres of 0.2 m and seven shell spheres of 0.25 m; three spheres of 0.32 m to 0.45 m on the arm and four of 0.4 m to 0.55 m on the boom), four spheres on the upper structure (three on the counterweight at the tail-swing radius and one at the cab roof), boxes of the cab, the deck, the boom foot and the track frames, and three footprint circles of radius 1.75 m along the undercarriage.

Seen from above, the spheres and the footprint circles are the machine's *collision circles*; workers are vertical cylinders, and their rows use the horizontal distance. While a loaded bucket is carried, both layers add a payload-retention row (Section 4.3). The admissible floor is the ground outside the excavation zone and, inside it, the floor of the current pass (Section 7) with a smooth batter beyond the zone's edges. Each row provides the gradient $J_i = \partial h_i/\partial x$ mapped to the coordinates and an explicit rate $b_i = \partial h_i/\partial t$, so that $\dot{h}_i = b_i + J_i u$. With the separation monitor active, workers enter the rows as static disks ($b_i = 0$), and a worker row whose standoff a person has entered no longer asks for the standoff to be restored ($\tilde{h}_i$ is replaced by $\max(\tilde{h}_i, 0)$ in both layers, so the condition becomes $\dot{h}_i \geq 0$): the rows keep the machine from moving toward a person but never command an evasive motion, and the approach of a person is handled by the monitor, which reduces speed or inhibits motion, the two forms of collision avoidance covered by ISO 21815-1 for earth-moving machinery [22]. The standoffs are 2° (joints), 0.25 m (self), 0.03 m (bucket vs. floor, so that the cutting edge can reach the design grade), 0.10 m (arm and boom vs. ground), 0.50 m (obstacles), 1.50 m (workers) and 0.40 m (footprint). Rows are grouped into the categories joint, self, ground, obstacle, worker and footprint. Here $h_i = 0$ is the modelled geometric boundary and $\tilde{h}_i = 0$ the requested clearance.

Closest-pair selection makes the distance rows piecewise smooth. The derivatives below apply locally while the selected pair is unchanged and the distance field is differentiable; re-selecting the rows at the next control step is not a proof of invariance at a pair switch, and an unobserved object supplies no row. These limits are the reason for the separate ground-truth clearance evaluation of Section 9.

4.2. Synchronized motion planner

The planner integrates the intents with jerk-limited second-order dynamics into a desired task-space motion (p_d, v_d, a_d), $p = (\psi, s, z, \phi)$, and couples its clock to the machine. With the weighted plan error $e = \|W(p_d - p(q))\| + \kappa_e \|(u_r - u) \oslash u_{\max}\|$, with the rate limits $u_{\max}$ of the reference layer (Section 4.3; $\oslash$ and $\odot$ denote element-wise division and multiplication), the time scaling

$$\rho = \operatorname{clip}\left(1 - \left(\frac{(e - e_{db})_+}{e_{\max}}\right)^2, 0, 1\right), \qquad v^\star = \rho\, v \odot v_{\max}, \tag{1}$$

slows the plan when the machine falls behind ($v^\star$ is the velocity command of the jerk-limited integrator), and the re-anchoring

$$\dot{p}_d = v_d - k_s \frac{(e - e_0)_+}{e}\,(p_d - p) \tag{2}$$

pulls the plan toward the machine when the error exceeds e_0 (the correction is zero at $e = 0$). The nominal joint rate solves a weighted least-norm problem [47], $\min \|W_t(J_p\dot{q} - \bar{v})\|^2 + \sum_j \mu_j \dot{q}_j^2$ with the Jacobian J_p of p with respect to the joint angles, a task-space weight W_t, $\bar{v} = v_d + K_p(p_d - p)$ and penalties μ_j that grow inside a 12° band before a joint limit, so that the motion is redistributed to the other joints before the joint-limit rows must act. Finally, the plan is scaled to the pump capability: with the steady flow demand $Q(u_n)$ (cylinder areas times moment arms, swing and travel displacements; for travel $Q_{\mathrm{tr}} = D_{\mathrm{tr}}(|v - B\omega/2| + |v + B\omega/2|)$ with the travel-motor flow per unit track speed D_{tr} and the track gauge $B = 2.6\,\mathrm{m}$) and the capability $Q_{\mathrm{cap}} = 0.85 \min(q_{\mathrm{del}}, P_{\max}/p_{\mathrm{ref}})$, with the pump delivery q_{del} at the engine speed, the available pump power $P_{\max}$ and a reference working pressure p_{ref}, $u_n \leftarrow \min(1, Q_{\mathrm{cap}}/Q(u_n))\, u_n$, and the planner state is scaled with it; otherwise the pumps saturate, the most loaded function can stall and the realized motion can leave the planned direction. In the same way, the plan is scaled to the admissible collision-circle speed v_{adm} of the separation monitor (Section 6.2): u_n and $\dot{u}_n$ are multiplied by the largest factor $\varsigma \leq 1$, found by bisection, for which no collision circle moves faster than v_{adm}, neither with the nominal motion nor with the nominal next reference before filtering. The latter, $u_r + a_n \Delta t$ with the nominal acceleration $a_n = \dot{u}_n + \kappa_r(u_n - u_r)$ of Section 4.3, leads the nominal motion by the planned acceleration at the onset of a motion. With the horizontal Jacobians $J_{c,j}$ of the circles, the two conditions are $\max_j \|J_{c,j}\, \varsigma u_n\| \leq v_{\mathrm{adm}}$ and $\max_j \|J_{c,j}(u_r + (\varsigma\dot{u}_n + \kappa_r(\varsigma u_n - u_r))\Delta t)\| \leq v_{\mathrm{adm}}$.

Proposition 1 (No accumulation of the plan error)**.** *Let $\bar{V} \geq \|W(v_d - \dot{p})\|$ bound the mismatch between the plan velocity and the machine's task velocity, which is finite because v_d and the actuator rates are bounded. For the continuous-time re-anchoring law in a consistent local angle coordinate, the geometric plan error $e_g = \|W(p_d - p)\|$ satisfies $e_g(t) \leq \max\{e_g(0),\, e_0 + \bar{V}/k_s\}$ for all $t \geq 0$.*

Proof. $\frac{d}{dt}\frac{1}{2}e_g^2 = (p_d - p)^\top W^2(v_d - \dot{p}) - k_s\frac{(e-e_0)_+}{e}e_g^2 \le e_g\bar{V} - k_s(e_g - e_0)_+e_g$, using $e \ge e_g$ and the monotonicity of $(e - e_0)_+/e$ in e. Hence $\dot{e}_g < 0$ whenever $e_g > e_0 + \bar{V}/k_s$. □

This bounds the internal plan state, not obstacle clearance or task completion. The same bound holds for the explicit Euler update when $0 \le k_s\Delta t \le 1$ and $\bar{V}$ bounds the step-averaged mismatch. Without re-anchoring (plan-then-track), the error grows for as long as the machine is blocked, e.g., by a barrier or a stall, and the subsequent catch-up produces large corrective references. The values used are $e_{db} = 0.3$, $e_{\max} = 1.2$, $e_0 = 0.6$, $k_s = 1.5\,\mathrm{s}^{-1}$, $K_p = 2\,\mathrm{s}^{-1}$, $W = \mathrm{diag}(1/0.5\,\mathrm{rad}, 1/0.3\,\mathrm{m}, 1/0.3\,\mathrm{m}, 1/0.3\,\mathrm{rad})$, so that e is dimensionless, and $\kappa_e = 0.25$.

4.3. Reference layer

Following [17], the reference layer filters the acceleration of the reference rather than its velocity, so that the reference remains trackable. With the reference slack $z_i = b_i + J_i u_r + \alpha_{1,i}\tilde{h}_i$, the measured barrier rate $\dot{h}_i = b_i + J_i u$ and the nominal acceleration $a_n = \dot{u}_n + \kappa_r(u_n - u_r)$,

$$
\begin{aligned}
a_r = \arg\min_{a}\ & \tfrac{1}{2}\|a - a_n\|^2_{Q_a} \\
\text{s.t.}\ \ & J_i a \ge -\big(\dot{J}_i u_r + \alpha_{1,i}\dot{h}_i + \beta_i z_i\big)\ \ \forall i, \\
& |a| \le a_{\max}, \qquad |u_r + a\Delta t| \le u_{\max},
\end{aligned}
\tag{3}
$$

and $u_r \leftarrow u_r + a_r\Delta t$. With b_i held constant locally, the constraint enforces $\dot{z}_i \ge -\beta_i z_i$ in the continuous-time approximation, so that $z_i(t) \ge z_i(0)e^{-\beta_i t}$ for a smooth row with feasible, zero-slack constraints. Because z_i concerns the reference rate u_r and the physical barrier derivative the measured rate u, reference consistency alone does not establish invariance of the physical machine. A further row, scaled so that it yields to the barrier rows where they conflict, keeps the cylinder flow of the next reference rate within the pump's capability. Constraints that cannot be met are relaxed with a slack per row, and such steps are counted as reference-layer infeasibility. We use $\alpha_1 = (1.5, 1.0, 1.0, 0.8, 0.8, 0.8)\,\mathrm{s}^{-1}$ per category, $\beta = 6\,\mathrm{s}^{-1}$, $\kappa_r = 8\,\mathrm{s}^{-1}$, and for (swing, boom, arm, bucket, travel speed, yaw rate) $u_{\max} = (0.6, 0.35, 0.6, 1.0)\,\mathrm{rad/s}$, $0.6\,\mathrm{m/s}$, $0.3\,\mathrm{rad/s}$ and $a_{\max} = (0.6, 1.0, 1.5, 3.0)\,\mathrm{rad/s^2}$, $0.4\,\mathrm{m/s^2}$, $0.3\,\mathrm{rad/s^2}$. The weights Q_a (and H below) are the identity, except that the bucket coordinate is weighted 20 times while a loaded bucket is carried, so that a correction is made with swing, boom and arm rather than by opening the bucket, and a payload-retention row keeps the bucket pitch of a loaded bucket within the range that holds the load. The yaw rate is weighted 10 times during travel, so that the filters brake rather than steer the machine off its path.

4.4. Actuator layer for valve-controlled actuators

The actuator layer acts on the duties. Each function f is characterized by a dead time $T_{d,f}$, a first-order lag, a deadzone and a steady duty-to-rate map $V_f(d)$, identified on the twin (dead times of 0.10 s to 0.15 s and lags of 0.05 s to 0.15 s, with 0.2 s and 0.55 s for the swing); $w = V(d)$ denotes the commanded steady coordinate rate. The actuator controller proposes the duties d_π by inverting the maps at the reference rate advanced by the lag, $u_r + \tau a_r$, with the duty scaled linearly to zero below a small rate band (the deadzone-aware inverse valve model), and adds a proportional feedback of the rate error $u_r - u$. Commands issued now take effect after their identified delays; the local predictor approximates the in-flight rate by $w^- = V(d_{k-1})$, the rate of the last applied duty (two to four commands are in flight during a dead time; they are assumed close to the last one, which the acceleration-error margin below covers only approximately).

Interval prediction. Under a first-order lag, each coordinate rate relaxes monotonically from the measured rate u_j toward w_j^-; a coordinate that accelerates beyond the nominal map (e.g., an overrunning, gravity-assisted load) is extrapolated with its measured trend $\dot{u}_j$.

Proposition 2 (Common-horizon dead-time bound)**.** *For a static row ($b_i = 0$), fix a prediction horizon $T_p > 0$. Assume that over $[t_k, t_k + T_p]$ (a) h_i is affine in the coordinates with constant gradient J_i, and (b) each rate $u_j(t_k + s)$ lies in the interval spanned by $\{u_j(t_k), w_j^-, u_j(t_k) + T_p\dot{u}_j(t_k)\}$ enlarged in each direction by $a_{\mathrm{err},j}s$, $0 \le s \le T_p$, with an acceleration-error bound $a_{\mathrm{err},j}$. Then $\tilde{h}_i(t_k + T_p) \ge \hat{h}_i(T_p)$, where*

$$
\hat{h}_i(T_p) = \tilde{h}_i + T_p\sum_j \min\big(J_{ij}u_j, J_{ij}w_j^-, J_{ij}(u_j + T_p\dot{u}_j)\big) - \tfrac{1}{2}T_p^2\sum_j |J_{ij}|a_{\mathrm{err},j}. \tag{4}
$$

Proof. Integrate $\dot{h}_i = \sum_j J_{ij} u_j$ over the interval. Each contribution is at least its minimum interval value minus $|J_{ij}| a_{\mathrm{err},j} s$; integration yields (4). □

The deployed predictor uses $a_{\mathrm{err},j} = 0.1\, a_{\max,j}$ and the separately identified delays $T_{d,j}$ in place of the common horizon, i.e., $\hat{h}_i = \tilde{h}_i + \sum_j T_{d,j} \min(\cdot) - \frac{1}{2} \sum_j T_{d,j}^2 |J_{ij}| a_{\mathrm{err},j}$; with unequal delays this is a local delay-compensation estimate rather than a bound at one common future time.

Effort-level condition. As in [17], the velocity-level barrier condition is $\Psi_{1,i} = \dot{h}_i + \alpha_{1,i} \tilde{h}_i \geq 0$. At the valve level we impose $\Psi_{2,i} = \dot{\Psi}_{1,i} + \alpha_{2,i} \Psi_{1,i} \geq 0$ and model the closed-loop rate response as $\dot{u} = (w - u)/\tau_c$. Here τ_c is the closing time constant of this condition, a design parameter of the filter rather than the response time of the hydraulics, whose dead time and lag the predicted barrier value $\hat{h}_i$ accounts for. Holding b_i constant gives the condition, linear in w and evaluated at the predicted barrier value,

$$J_i \frac{w - u}{\tau_c} + \dot{J}_i u + (\alpha_{1,i} + \alpha_{2,i})(b_i + J_i u) + \alpha_{1,i} \alpha_{2,i} \hat{h}_i \geq 0. \tag{5}$$

Proposition 3 (At-rest compatibility of the layers)**.** *For a static row with $\tilde{h}_i \geq 0$ and condition* (5) *without the braking credit introduced below, take $\alpha_{2,i} = 1/\tau_c$. At rest ($u = 0$, $\dot{u} = 0$, $w^- = 0$, $b_i = 0$), with the acceleration-error margin of* (4) *neglected, condition* (5) *reduces to $J_i w \geq -\alpha_{1,i} \tilde{h}_i$, the velocity-level inequality $z_i \geq 0$ of the reference layer. A reference that satisfies it is therefore admitted by condition* (5) *at rest; for $\alpha_{2,i} < 1/\tau_c$ the condition admits only the fraction $\tau_c \alpha_{2,i}$ of the admissible approach rate.*

Proof. Multiply (5) by τ_c and substitute $u = 0$, $b_i = 0$, $\hat{h}_i = \tilde{h}_i$: $J_i w + \tau_c \alpha_{1,i} \alpha_{2,i} \tilde{h}_i \geq 0$. □

A mismatch would matter in practice: a command that the actuator layer reduces near a barrier can fall into the valve deadzone, so that the machine stops short of the barrier (lock-in). The braking credit introduced below makes the deployed condition stricter: at rest it becomes $\sum_j \min(J_{ij} w_j, 0) \geq -\alpha_{1,i} \tilde{h}_i$, which admits a reference whose approaching coordinates stay within the margin but does not count the retreat of one coordinate against the approach of another, as the reference layer does. The proposition therefore does not hold for the deployed filter when a reference combines approaching and retreating coordinates: such a reference is reduced at rest, and the lock-in described above can occur; this is the price of not relying on corrective commands. We use $\tau_c = 0.05\,\mathrm{s}$, i.e., $\alpha_2 = 20\,\mathrm{s}^{-1}$. The proposition concerns the two layers only: a reference rate that both admit can still be so small that the inverse valve map commands it below the valve opening, so that the machine approaches a barrier slowly or stops short of it; this limits the penetration at the pass start (Section 7.2).

Braking credit. Condition (5) credits every corrective command in full. On a valve-controlled machine a small corrective command may fall into the deadzone, the pumps may not supply the flow, or a loaded function may stall, whereas closing a valve is the more dependable way to arrest flow. Each coordinate's contribution is therefore credited only as far as it reduces motion toward the barrier, or up to the measured rate of a coordinate already moving away. Writing (5) as $J_i w \geq r_i$ with $r_i = c_1 J_i u - \tau_c(\dot{J}_i u + (\alpha_1 + \alpha_2) b_i + \alpha_1 \alpha_2 \hat{h}_i)$ and $c_1 = 1 - \tau_c(\alpha_1 + \alpha_2)$, the credit-limited condition is

$$g_i(w) = \sum_j \min(J_{ij} w_j, c_{ij}) - r_i \geq 0, \qquad c_{ij} = \max(J_{ij} u_j, 0). \tag{6}$$

Filter. The actuator layer solves

$$w^\star = \arg\min_w \ \|w - V(d_\pi)\|_H^2 \quad \text{s.t.} \quad g_i(w) \geq 0 \ \forall i, \quad \underline{w} \leq w \leq \overline{w}, \tag{7}$$

with componentwise limits from the identified maps, and recovers the duties with the deadzone-aware inverse map. Each g_i is concave, hence the feasible set is convex: it is the intersection of the 2^n half-spaces ($n = 6$ coordinates) $\sum_{j \in S} J_{ij} w_j + \sum_{j \notin S} c_{ij} \geq r_i$ over all subsets S of coordinates, each of which is necessary. We solve (7) by cutting planes: starting from the subset that is tight at the proposed command, the subset tight at the current solution is added for every violated row and the quadratic program (QP) is re-solved with the dual active-set method of [48].

Proposition 4 (Exact solution)**.** *If $\boldsymbol{H}$ is positive definite,* (7) *is feasible and every relaxed QP is solved exactly, the uncapped cutting-plane iteration terminates after finitely many iterations with the unique solution of* (7)*.*

Proof. Every cut is necessary, so each relaxed feasible set contains the exact one and the relaxed solution is optimal whenever it is feasible. If it violates row i, the cut of the subset tight at it is violated as well and is not yet in the relaxation; there are finitely many subsets. □

The implementation caps the iteration at six cuts; subproblems with remaining violations are relaxed with a slack per active inequality (quadratic penalty 10^4), and a command that remains infeasible is replaced by braking (all functions to neutral) when braking has the larger worst-row margin. At rest within the dead-time margin of a barrier ($r_i > 0$), condition (6) credits no corrective command, since no coordinate moves away yet: every command is infeasible, braking is selected at every step, and the machine stays at rest indefinitely, for instance with the bucket between a ground row and the standoff of an obstacle row, which boom, arm and bucket move in opposite directions. In the missions, the single pits and the worker study, a machine at rest (every measured rate below $0.01\,\mathrm{rad/s}$, or $0.01\,\mathrm{m/s}$ for the travel coordinates) with an instance that remains infeasible is therefore filtered once more, with the condition relaxed: an infeasible row only must not decrease ($r_i \to 0$), every other row keeps its requirement, and the ground rows, which keep the front on the soil side of its admissible floor and protect neither people nor the machine, may decrease at most at $0.2\,\mathrm{m/s}$. Braking meets the relaxed condition, so the relaxed QP is feasible; its solution is applied if it is not zero, except during a protective or commanded stop. Once the machine moves away from the barriers, its measured motion is credited and condition (6) is feasible again; a solution whose duties lie inside the valve deadzone, however, leaves the machine at rest (a case that the task layer resolves for the loaded swing, Section 7.3). The benchmarks and the learning studies on the reduced-order twin and the stopping-bound identification use the layer without this rule (it does not act during a stop). All invariance interpretations require a suitable initial state, an adequate geometry and actuator model, feasible constraints and a valid treatment of sampling at the $20\,\mathrm{Hz}$ control rate; the evaluation therefore measures raw boundary crossings and clearance deficits directly.

5. Perception

The perception stack turns the LiDARs and cameras into what the task layer, the safety layers and the separation monitor need: the terrain (dig progress and spoil heaps), the static obstacles (a distance field), the workers and the truck (with velocities), the positions of persons who are temporarily not observed, and the task semantics (where the pits and the spoil area are). It exposes the same interface as the exact world of the *privileged* experiments (exact soil, obstacles and actor states; Section 9), so that all layers run unchanged on perceived information.

5.1. LiDAR mapping

The mapping follows the layered design of nvblox [39], re-implemented on the CPU. Full scans (every $2\,\mathrm{s}$) are integrated into a truncated signed distance field (TSDF; $0.15\,\mathrm{m}$ voxels, truncation $0.45\,\mathrm{m}$, weight cap 30) with free-space carving along the rays, so that voxels of objects that have moved away are cleared. Ground returns update a 2.5-D elevation layer ($0.1\,\mathrm{m}$ cells, exponential averaging). Where the cab-roof sensor cannot see the floor of the pit being dug, the elevation layer is lowered along the path of the bucket's cutting edge (the line through the teeth, extended to the bucket width, every $0.05\,\mathrm{m}$), since the terrain there is not higher than the edge; later returns raise it again where soil flowed back. A Euclidean signed distance field (ESDF) is computed at $0.25\,\mathrm{m}$ resolution from the static surface voxels more than $0.35\,\mathrm{m}$ above the terrain, without the tracked dynamic objects, in a $40\,\mathrm{m} \times 40\,\mathrm{m}$ window around the machine; the obstacle rows query it with trilinear interpolation and its gradient. Within a confirmed pit and $1.3\,\mathrm{m}$ around it (the pit's soil zone, which covers the stepped walls of the pit, the notches that the passes cut beyond its near wall, Section 7, and the soil that the bucket pushes and drops at their ends), surface voxels less than $0.3\,\mathrm{m}$ above the original ground are terrain. Soil heaped there rises above the elevation estimate before the LiDAR has updated it; taken for an obstacle, it would keep the bucket at the obstacle standoff, and with a narrower zone ($0.55\,\mathrm{m}$) it stopped drags short of the near wall and pushed the closing bucket back into the pit in preliminary runs. Sparse rays

do not clear every voxel that a walking worker or the departing truck occupied; a 2-D grid (0.25 m cells) therefore records when a tracked moving person (radius plus 0.3 m) or the moving truck last covered a cell, and surface voxels there are excluded from the static layer unless they have been hit again since. Objects taller than 0.65 m are extruded to the ground.

5.2. Dynamic objects

Object scans (5 Hz) without ground returns are segmented in the range image of each LiDAR in the manner of [40], so that objects that are close in the map but separated in range (a worker in front of a container) are not merged, and tracked with constant-velocity Kalman filters. A track is a worker if its extent and height are person-like, it has been confirmed in at least three scans and it has moved: its estimate lies more than 0.6 m from where it was first seen, at a speed below 2.5 m/s, while less than half of its segment's points lie on long-observed static surface (TSDF weight at least 8). A person-sized part of a structure can pass this test, since the centroid of its segment moves with the view, as parts of the structure are hidden or come into view, before its voxels are long-observed; the status would then hold for as long as the structure stays in view. The status is therefore revoked when the track's own segment (diagonal at most 1.2 m, so not a person merged with a structure) has had at least 90 % of its points on long-observed static surface for 10 s while the track moved slower than 0.5 m/s; a revoked track becomes a worker again only by moving 0.6 m once more. A worker who stops keeps the status, because the voxels of a confirmed walking person that become long-observed are marked as that person's surface and do not count as static (below). The radius of person tracks is capped at 0.6 m, and a person associated with a merged segment is measured from the points within 0.6 m of the prediction. A large moving or parked segment of truck size is the truck. The fraction of a segment's points in TSDF voxels with a high integration weight distinguishes static structure (e.g., the fence) from moving objects. The truck's pose follows from its high returns (bed walls and cab above 2.3 m), and its load from the lower quartile of the returns inside the walls; a parked truck keeps its last pose while it is occluded by the boom, and a truck that was moving when it was last observed is removed after 5 s without observation.

5.3. Workers lost from view

A worker hidden by the truck, the boom or a pit edge must not vanish from the safety layers. For every worker track with at least five associations, at least two of them with a person-like shape, that is lost within 16 m of the swing axis, the perception keeps the set of positions the person could occupy. The set lives on a 0.25 m grid, grows every object scan by the walking speed $v_h = 1.6\,\mathrm{m/s}$ and loses every cell that is observed free at person height (0.3 m to 1.6 m above the terrain, up to the first return), as well as the cells under the undercarriage and inside the perceived truck body. As long as rays through a person's body would have returned, the set contains every position of a person who walks no faster than v_h. A person track that reappears inside a set takes it over and is reported at once. A set that leaves the 16 m region around the swing axis is dropped, since the near-field observations no longer prune it there and it would soon cover the surroundings of the machine; a person who walks back toward the machine while still hidden is therefore not covered (condition (E2) of Section 6.3). Two rules keep the free-space carving conservative and complete:

- A ray that passes beside a standing person, or between the legs, would mark the person's own cell as free. Cells within 0.45 m of object returns that lie neither on long-observed static surfaces (TSDF weight at least 8) nor inside the perceived truck body are therefore never cleared.
- The two LiDARs (on masts 4.05 m above the ground on the cab roof and 3.65 m on the deck) do not see the ground-level ring around the tracks below the roof line of the upper structure. A set that grows into this ring could not be cleared by them, and a monitor that respects the set would stop the machine until the set expires after 600 s. Four near-field scanners on the corners of the undercarriage, as used for person detection around mobile machinery, observe this ring: each has one horizontal scan plane 0.5 m above the ground, a 270° field of view facing outward, 0.5° resolution and 10 m range. A scanner ray observes the cells along it up to its first object return (less 0.3 m) while the plane lies within the person band above the elevation map; beyond a rise of the ground toward the plane (a spoil heap) the rest of the ray is not used.

Persons near the machine. Close to the machine, the legs of a person are below the lower edge of the mast LiDARs' field of view (within about 3 m of the cab-roof unit and 6 m of the deck unit, from their mounting heights and field-of-view limits), so a person there is not classified from them, and a person who has stood still since the machine started is never classified by the tracker, whose person class needs motion. The near-field scanners therefore also detect persons, as the protective fields of safety laser scanners on mobile machinery do [49]: every cluster of their returns (single linkage at 0.5 m, at least two returns) whose diagonal is at most 1.2 m and that lies neither on long-observed static surfaces nor on the truck (within 0.6 m of the truck's mast-LiDAR returns of the same object scan, or in its perceived body, inflated by 0.5 m for its pose error and, while the truck moves, by its travel in 0.5 s) is a person for the separation monitor, whether or not it moves. Beyond 8 m from the swing axis, where the mast LiDARs observe a person's legs, a return also counts as static structure when long-observed static surface lies above it within 0.3 m horizontally (up to 2.5 m above the terrain): the face of an object that points away from the mast LiDARs is seen only by a scanner, whereas its top is long-observed, as for the end face of the pipe stack behind P8 seen from the stance of P8. A person who stops is integrated by the TSDF in the same way as static structure; voxels that become long-observed surface within 0.6 m of a confirmed walking person are therefore marked as that person's surface and do not make returns static. A detection is held for 1 s after it was last seen, its radius growing with the walking speed. A confirmed walking person, in turn, keeps its class while only its upper body is observed.

The persons of the separation monitor (Section 6) are the tracked workers, the near-field detections and the cells of the sets, inflated by the distance $v_h t_o = 0.32\,\mathrm{m}$ walked between two object scans ($t_o = 0.2\,\mathrm{s}$); for the barrier rows, each set is represented by its cell nearest to the machine's collision circles.

5.4. Camera semantics

A compact encoder–decoder network (U-Net style [41], three scales with 16/32/48 channels, output at half resolution) labels every pixel of the three cab cameras as background, dig-site marking (orange paint), dump-zone marking (blue paint and flags), truck, person, other obstacle or own machine. It was trained on 1,200 rendered images (validation on 200 others) with random machine poses, swing angles, front poses, worker and truck positions and partially dug pits, with labels from the segmentation rendering of the simulator; on the validation images, the intersection over union is 0.62 for the dig-site marking, 0.64 for the dump-zone marking, 0.85 for the truck and 0.58 for persons (0.78 averaged over the classes). At run time one camera image is processed every 0.5 s, the three cameras in turn (every 4 s during excavation, while the outline of the pit is fixed); labelled pixels within 15 m are back-projected onto the elevation map and accumulated in a bird's-eye-view evidence grid, from which dig sites and the dump zone are extracted by morphological closing, hole filling and minimum-area rectangles and matched to the work order (Section 7.1).

6. Worker Protection

The worker rows of the barrier layers keep the machine from moving toward a person, but they cannot keep a standoff against a person who walks toward a machine whose swing needs seconds to stop. After a stop command during a fast loaded swing, the front of the twin still travels up to 2.91 m (Section 10.6). Worker protection therefore follows the logic of speed and separation monitoring: the machine moves only as fast as it can stop before any person, approaching at walking speed, could come closer than the requested standoff. This section describes the stopping bounds on which the monitor rests, the monitor itself and the operating envelope within which its design rule is claimed.

6.1. Identified stopping bounds

The monitor needs two properties of the machine's controlled stop (all intents zero, planner held at the measured state, both safety layers active): how far each collision circle can still travel after the stop command, and how long the machine can still move. Both were measured on the full twin from stops commanded at random times during coverage excavation with truck loading and spoil dumping (digging, loaded and empty swings, dumping). For every stop, the circles of the front, the upper structure and the footprint were followed for 6 s; the travel d_j of circle j is its largest displacement from its position at the command, and the stopping time t_s is the time after which no circle moves by more than 0.05 m any more (it includes the swing's rebound

and the decay of the hydraulic oscillation). A stop commanded while the machine accelerates travels farther and lasts longer than one commanded at the same speed at a steady motion, because the valves keep the motion going during the stopping lag; at the onset of a motion the valves are open before the circles move, and the machine still approaches the speed of its motion reference after the command. With v the largest circle speed at the command, a the largest rate at which a circle speeds up (its acceleration along its velocity, from the circle positions of the last control steps) and v_r the largest circle speed of the motion reference (the reference rates u_r mapped through the Jacobians of the circles), the bounds are therefore evaluated at the effective speed

$$v_e = \max\left(v + \max(a - a_{\mathrm{acc}}, 0)\, t_l,\ v_r\right), \tag{8}$$

as

$$d(v_e) = v' t_l + \frac{v'^2}{2a_s} + d_0, \quad v' = v_e + a_{\mathrm{acc}} t_l, \qquad T(v_e) = \min\left(T_a + k_T v_e,\, T_{\max}\right), \tag{9}$$

with the stopping lag $t_l = 0.25\,\mathrm{s}$, the deceleration $a_s = 2.5\,\mathrm{m/s^2}$ and the acceleration $a_{\mathrm{acc}} = 1.5\,\mathrm{m/s^2}$ that the bound covers without the correction of Eq. (8). The time bound covers every stop with a margin of 0.4 s (two object-scan periods): T_a is the longest stopping time of the stops commanded at rest ($v_e < 0.05\,\mathrm{m/s}$; zero if there is none) plus the margin, $T_{\max}$ the longest stopping time plus the margin, and k_T the smallest slope that covers the stops in between; d_0 is the smallest multiple of 0.05 m with which every stop is covered. A machine whose circles and reference are at rest has nothing to stop, and a motion that is starting is evaluated at the speed of its reference (Eq. (8)); T_a therefore only has to cover a machine at rest, and it decides how close a person may be to such a machine. No measured stop was commanded at rest (Section 10.6), so T_a is the margin alone, set by this argument rather than by a measurement. For a moving stop ($s_{\max} > 0$), the travel of every circle must be covered by its share of the fastest circle's bound, $d_j \le 1.1\, d(v_e)\, s_j / s_{\max} + d(0)$, where s_j is the circle's speed at the command in the faster of the measured and reference motions and $s_{\max} = \max_j s_j$.

The bounds were identified on 60 stops (P3 with scenario seeds 101 and 102 and P8 with seed 101, 20 stops each) and checked on 40 further stops (P3 with seed 103 and P8 with seed 102). The bounds used in all experiments apply the same rule to the union of both sets, which gives the identified values, $T_a = 0.40\,\mathrm{s}$, $k_T = 2.46\,\mathrm{s^2/m}$, $T_{\max} = 2.65\,\mathrm{s}$ and $d_0 = 0.30\,\mathrm{m}$, and they were checked on a third set of 60 stops, measured after they were fixed (P3 with seed 104 and P8 with seeds 103 and 104; Section 10.6). They are empirical envelopes of the measured stops of the twin during excavation, not worst-case bounds; stops during travel were not identified. They are tested in operation by the worker-protection experiment of Section 10.6 (P3 with scenario seed 301 and workers on prescribed paths), which measures the clearance that the machine's own motion leaves to a person.

6.2. Separation monitor on the swept stopping region

The protective separation distance of speed and separation monitoring [20] adds, to the requested standoff, the distance a person walks during the reaction and stopping time of the machine and the distance the machine travels in that time. Treating the machine's travel as a disk around each collision circle is conservative on an excavator, whose fastest motion is the swing: during a stop, the circles move on arcs about the swing axis, and a person beside the arc cannot be reached by the stop. For every circle j (centre p_j, radius r_j, velocity v_j, speed s_j), the monitor splits the velocity of the faster of the measured and the reference motion (the reference at the onset of a motion) into the part due to the swing rate, $v_{\mathrm{sw},j} = \dot{\psi}\, e_z \times (p_j - p_0)$ (e_z the vertical unit vector) about the swing axis p_0, and the rest, $v_{\mathrm{res},j} = v_j - v_{\mathrm{sw},j}$ (boom, arm, bucket, travel). For $s_{\max} > 0$, define $A = 1.1\left(t_r + d(v_e)/s_{\max}\right)$. The circle sweeps an arc of length $A\|v_{\mathrm{sw},j}\|$ about the swing axis in the direction of the swing (and 0.3 m back, for the rebound), inflated by a disk of radius $A\|v_{\mathrm{res},j}\| + d(0)$. A person k with radius r_k at distance D_{jk} from this swept region must satisfy

$$\delta = \min_{j,k}\left(D_{jk} - r_j - r_k\right) - v_h\left(t_r + T(v_e)\right) - \sigma \ \ge\ 0, \tag{10}$$

with the walking speed $v_h = 1.6\,\mathrm{m/s}$ of ISO 13855 [21], the detection period $t_r = 0.25\,\mathrm{s}$ (object scan and control step) and the worker standoff $\sigma = 1.5\,\mathrm{m}$. For a non-swinging machine at a steady motion, the swept region

reduces to the disk form of the protective separation distance, $S_p(v) = \sigma + v_h(t_r + T(v)) + 1.1\,(vt_r + d(v)) + d(0)$. When both the measured and reference circle speeds are zero, the monitor evaluates clearance against $S_p(0)$. The persons are the tracked workers, the person detections of the near-field scanners and, for every worker lost from view, the cells of its reachable set (Section 5), inflated by the distance walked between two object scans.

While $\delta < 0$ at the current motion, the machine performs the controlled stop; the stop is released once the machine is nearly at rest ($\max_j s_j < 0.1\,\mathrm{m/s}$) and $\delta \geq 0.3\,\mathrm{m}$ for a machine at rest. Otherwise, the monitor limits the speed of the motion that the planner commands: the planner's nominal rates are time-scaled, like for the pump-flow capability (Section 4), so that no collision circle of the nominal motion moves faster than an admissible speed v_{adm}, computed with the circles' Jacobians at the current pose. A direction-free cap follows from the extent of the swept condition for a straight motion: a circle moving at the largest speed v straight toward a person needs the separation $\sigma + v_h(t_r + T(v)) + 1.1\,(vt_r + d(v)) + d(0)$, and with the clearance c to the nearest person, the circles may move at the speed $v_{\mathrm{allow}}(c)$ at which this separation equals $c - 0.3\,\mathrm{m}$. With v the largest measured circle speed, the admissible speed is

$$v_{\mathrm{adm}} = \begin{cases} 0.9\,\kappa\, v_{\mathrm{allow}}(c), & v < 0.3\,\mathrm{m/s}, \\ \max\left(0.9\,\kappa\, v_{\mathrm{allow}}(c),\ 0.9\,\lambda v\right), & \text{otherwise}, \end{cases} \tag{11}$$

with $\lambda \in [0, 1.5]$ the largest uniform speed factor of the current motion for which $\delta \geq 0.3\,\mathrm{m}$ (bisection). Below 0.3 m/s, in particular after a release, the direction of the coming motion is not yet known and only the direction-free cap applies; a faster machine keeps its direction, so that a swing away from a person is not slowed and a motion toward a person is slowed while it still keeps the separation of the release condition, before it would have to be stopped. Because the hydraulic motion overshoots its reference at the onset of a motion, the factor $\kappa \in [0.1, 1]$ corrects the cap with the speed $\bar{v}$, the larger of v and the largest circle speed of the reference: in every control step that follows one governed by the direction-free cap, $\kappa \leftarrow \max\left(0.1, \kappa \max(0.5, 0.9\,v_{\mathrm{allow}}(c)/\bar{v})\right)$ while $\bar{v} > 0.9\,v_{\mathrm{allow}}(c)$, $\kappa \leftarrow \min(1, 1.02\,\kappa)$ while the cap binds and $\bar{v} < 0.8 \times 0.9\,v_{\mathrm{allow}}(c)$, and κ is held otherwise (in particular while the factor λ governs and during a stop). In either case the swept condition is evaluated in every control step, and a motion that violates it ($\delta < 0$) is stopped. The monitor never commands a motion of its own: in line with ISO 21815-1 [22], collision avoidance is by speed reduction or motion inhibit, not by evasive manoeuvres. It is a design rule of this simulation study, not a certified safety function.

6.3. Operating envelope

The design rule claimed for the monitor is: *the machine's own motion never reduces the horizontal clearance between a person and the machine's collision geometry below the standoff σ, and no three-dimensional contact between a person and the machine occurs while the machine moves toward the person.* It is claimed only within the following envelope, whose conditions are tested separately in Section 10.6 (the machine's own motion is an approach toward the person faster than 0.05 m/s, Section 9; the slower creep of a stopped machine is reported with the results):

(E1) persons move at most at walking speed, $v_h \leq 1.6\,\mathrm{m/s}$;

(E2) every person within 16 m of the swing axis is tracked, covered by the reachable set of a track that was lost within this region, or detected by the near-field scanners (within their 10 m range, at leg height, not hidden behind another object and, beyond 8 m from the swing axis, not below long-observed static structure); since a person who stands still is classified only in the near field, a person must enter the region in view of the mast LiDARs or of the near-field scanners;

(E3) the machine moves within the speed range and the operating conditions of the identified stops (excavation motions on level ground, the soil and payloads of the identification set, circle speeds up to 3.5 m/s); travel between the pits is protected by the same monitor but lies outside this condition, since no travel stops were identified;

(E4) perception reports a tracked person within the position error covered by the person radius, and object returns are not lost for longer than the reachable sets bridge.

A person who walks into a machine at rest, or under its raised front, reduces the clearance by the person's own motion; the monitor stops or holds the machine, but such worker-induced clearance losses are a matter for

site rules (exclusion zones) and are reported separately from machine-induced ones. Running persons violate (E1); they are evaluated as a test outside the envelope. The envelope concerns protection, not productivity: with the bounds of the experiments, a person restricts the machine within 3.4 m of a collision circle at rest, 8.0 m of a circle moving toward the person at 1 m/s and 12.2 m of a bucket that swings toward the person at 3.5 m/s, and productive work needs a site rule that keeps people farther away; the site workers of the missions keep 12.5 m from the machine, except while they leave the zone at the end of the walkable site (Section 3.3).

7. Construction-Task Layer

The work order lists the pits (design rectangle, depth below the original ground and dump place) and the site plan gives the truck lane and the marked spoil area. A state machine sequences survey, pit selection, folding of the front, travel to the stance and excavation (Fig. 3). This section describes how a pit is excavated to its design grade and how the material is delivered; the parameters are those of the experiments (pits of 1.6 m × 1.6 m, 0.6 m deep).

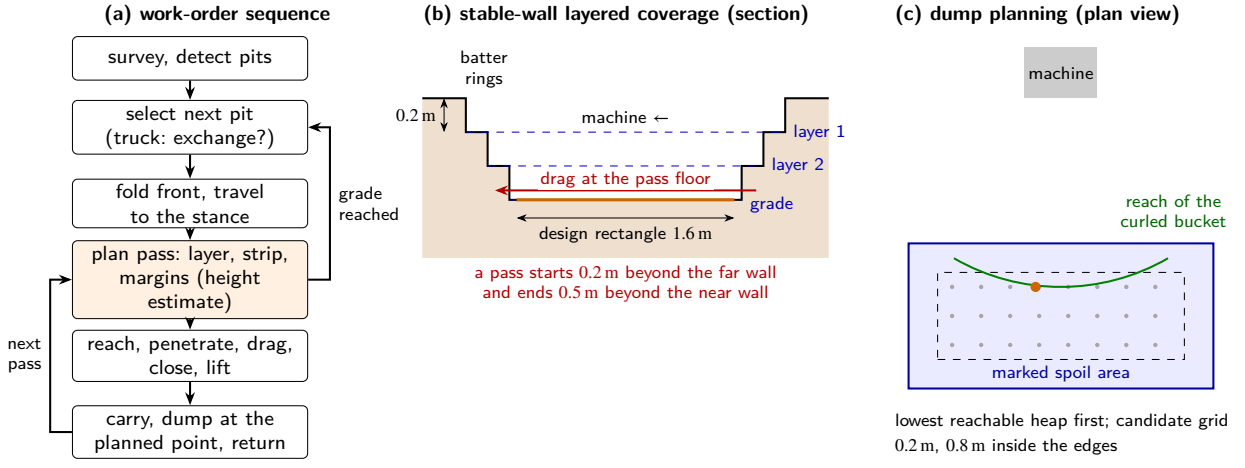

Figure 3: Construction-task layer. Left: the mission state machine. Middle: section through a pit with the three layers of the stable-wall coverage, the batter rings and the pass margins. Right: dump planning on the spoil area (reachable candidates of the curled bucket, lowest heap first).

7.1. Work order, grade reference and completion criterion

A pit is confirmed when the cameras find its orange marking within 0.6 m of the work-order position (a pit that is not found at its work-order slot is skipped and reported); from then on the georeferenced work-order rectangle is used, because the marking outline is located only to 0.1 m to 0.2 m, which would shift the boundary of a pit that is evaluated on a 0.1 m grid by one or two cells. The design grade is referenced to the site datum, and so is the task-space frame of the skills and the safety layers: the base of the twin moves on the datum plane, and a real machine would obtain the same reference from GNSS machine control. The reference therefore does not follow soil spilled under the machine. The evaluation counts a pit as complete only when every point of the design rectangle on a 0.1 m grid, the boundary lines included (289 points), lies within ±0.08 m of the design grade on the true soil. The tolerance is set by the twin: the soil is a 0.1 m heightfield, and the loose film that a drag leaves and the tooth radius take a few centimetres; tighter tolerances were not studied.

7.2. Stable-wall layered coverage

The soil model lets a wall of bank soil stand as long as the height step between neighbouring 0.1 m cells does not exceed $\tan(70°) \times 0.1\,\text{m} = 0.275\,\text{m}$; steeper walls fail, and the failed soil loosens and runs out at

the loose angle of repose (35°). A pit cut with vertical walls therefore refills its boundary every time a pass undercuts the wall toe. The coverage planner instead cuts the pit top-down in layers of 0.2 m. Its design surface is the grade over the rectangle and a working margin of one cell ring (0.1 m), and batter rings 0.2 m and 0.4 m above the grade outside it, with the original ground beyond. Every layer covers the cells whose design surface lies at or below its floor, so that the finished walls step by 0.2 m per ring (0.22 m at the floor of the grade layer, which the cutting edge cuts 0.02 m below the grade), below the failure step. The layer thickness also sizes the passes. The cycle closes the bucket when its content reaches 1.04 m^3 (85 % of the struck capacity of the bucket model, Section 10.1), and a bucket that closes before the drag has reached its planned end beyond the near wall drops soil into the pit. With 0.2 m layers, a pass of an upper layer carried 0.71 m^3 of loose soil on average in the ten-pit missions, and two of the 1000 carries started with more than the closing content (the largest load was 1.09 m^3). Layers of 0.3 m would exceed the failure step: their finished walls step by 0.32 m, and relaxing the finished profile of one pit in the soil model moves about 0.1 m^3 of wall soil onto the pit (none with 0.2 m layers; a unit test in the code archive repeats this computation); in a preliminary run with 0.3 m layers, the bucket content also passed the closing content over the near boundary row of the pit, where the closing bucket buried it. On an upper layer, the passes aim 0.05 m below the layer's floor (the lip rides 0.05 m to 0.1 m above its target while it cuts a thin layer at full reach), and the admissible floor of the ground rows (Section 4.1) lies a further 0.02 m lower over the layer's footprint, so that the barrier holds the teeth near the layer floor and its gradient does not push the bucket off the strip where the final floor rises. An upper layer is finished when its batter ring is cut (less than 0.05 m^3 above its floor); the deeper layer removes the remaining interior material. On the grade layer, the admissible floor lies 0.07 m below the grade: above it, the tooth radius (0.02 m) and the standoff of the bucket row (0.03 m) let the cutting edge reach the pass floor 0.02 m below the grade, so that the loose film that a drag leaves on the floor and isolated spill stay within the tolerance.

Each layer is split into strips that run parallel to the pit axis (the line from the swing axis through the pit centre), side by side, whose bucket footprints (1.376 m) cover the layer's footprint, with the bucket's edges at the outer cell edges. A pass penetrates at its far end and drags toward the machine at the pass floor; the cut target is this floor and does not depend on the surface estimate ahead of the lip, which in the perceived world still shows uncut cells at the strip edges. Two margins follow from the cutting process. The penetration stops 0.1 m to 0.15 m above the floor, where the descent rate that the ground row admits is so small that the inverse valve map commands it below the valve opening of the boom (Section 4.4), and the drag cuts the rest on its first decimetres; every pass therefore starts 0.2 m beyond the far wall. When the bucket closes at the end of a drag, the lip is held in the soil while the bucket rotates about it, so the bucket body sweeps away from the machine and drops part of the soil it carries. The cab-roof LiDAR does not see the floor behind the near wall, where such spill would go unnoticed; every drag therefore ends 0.5 m beyond the near wall, so that the closing bucket drops its soil beyond the rectangle (the thin passes below remove what slides back). The admissible floor of the ground rows is extended by the same margins on the far and the near side. Since every layer's passes extend beyond its own footprint, the notches that the margins cut beyond the near and far walls step by 0.2 m as well; the soil zone of the pit in the obstacle map covers them (Section 5).

A pass whose load exceeds 0.05 m^3 is followed by a carry and a dump; with a smaller load the next pass starts at once, and its reach keeps the bucket curled until the lip is over the middle of the pass, where the drag collects the load. Two alternatives failed in preliminary runs. A bucket opened at once drops such a load over the near rows, where the soil that falls out is hidden from the LiDAR (0.03 m^3 to 0.05 m^3 raised the near boundary row by up to 0.13 m). A bucket opened only 0.5 m before the pass start drops the load at the far boundary row, and the curled bucket keeps the lip so high that the penetration runs out of time above the floor, so that the drag passes over the dropped soil and leaves it as a ridge along that row (up to 0.12 m above the grade in preliminary runs of whole missions). On the grade layer, a held load is not collected reliably either: when the reach opens the bucket, the load falls out across the whole bucket width, and the drag of the next pass, often on the neighbouring strip, collects it only within its own cutting width, so that the rest stays on finished floor in the near rows (0.045 m^3 left one cell 0.13 m above the grade in a mission of the fifth earlier run, Section 9). On the grade layer, every load is therefore carried and dumped, however small, and the dump ends when less than 0.005 m^3 is left in the bucket, or after 2 s at the dump place when less than 0.05 m^3 is left (elsewhere at once below 0.05 m^3). The penetration ends when the teeth reach the

floor, when they stall for 1 s, or after 13 s. After every pass (at the dump, or at once when the load stays in the bucket), the next pass is planned from the height estimate of the pit: the true soil in the privileged world, and otherwise the LiDAR elevation map lowered along the path of the cutting edge. The layers are cut top-down once: the planner chooses the highest unfinished layer at or below the deepest layer it has reached, and its strip with the most material above the floor; on the grade layer, the strips that still cover a rectangle point outside the completion tolerance come first, since loose soil in the margin ring otherwise draws the passes to one side of the pit. A strip whose remaining material did not decrease in three passes in a row on the grade layer, or in two on an upper layer, is set aside (a pass counts as a decrease when it removes at least $0.01\,\mathrm{m}^3$ or a quarter of the strip's remaining material). A grade strip that still covers a rectangle point outside the tolerance is the exception: it is retried until the pit's time limit. A grade strip that was set aside returns to the plan when one of its rectangle points leaves the tolerance again. The far corners are the last points to reach the grade. A grade pass over a strip that is nearly at the grade—its estimated largest residual is below 0.2 m or its estimated remaining volume below $0.05\,\mathrm{m}^3$, since in the perceived map a few cells that neither the LiDAR nor the cutting edge has updated keep the largest residual high long after the strip is cut—is a *finishing pass*, and on the outer strips its start is shifted 0.04 m laterally away from the pit axis.

Every grade strip ends with a thin pass. A pass that removes a full layer nearly fills the bucket, and a short pass that collects a heap at the near end also carries a load; when the bucket closes beyond the near wall, part of this soil spills and slides back into the near rows. The cab-roof LiDAR does not see these rows behind the near wall, so the carved map still shows them at the grade. A pass therefore ends the work on its strip only if the strip's largest estimated residual was below 0.1 m when it started; otherwise the strip receives one more pass over its whole length, which carries little soil (at most two such passes per strip, since cells of the perceived map that are not observed again can keep the estimate high). The planner ends the pit when its height estimate lies at most 0.08 m above the grade at every rectangle point and every grade strip has ended with such a pass, when the remaining material of no strip is reduced any more, or at the time limit; whether the pit is complete is then evaluated on the true soil. The planner takes the decision that a pit is finished, by its estimate or because no strip improves, only at the end of a dump: when it falls due at the end of a pass, with the front still over the pit, the bucket is carried to the dump place first, and the planner decides on the estimate that the LiDAR has updated during the swing. The cutting edge lowers the estimate to its own height, so that soil that comes back behind the teeth at the end of a pass would otherwise not be seen before the decision.

7.3. Execution of the passes

The dig–dump cycle (reach, penetrate, drag, close, lift, carry, dump, return) is a scripted skill in task space (swing angle, radius and height of the bucket lip, and bucket pitch, the angle of the line from the bucket pivot to the teeth against the horizontal: $-90°$ with the teeth pointing straight down, more negative values curling the bucket toward the machine) with proportional intents and sub-phase time limits. Three properties set the lateral accuracy of a pass. First, the inverse map from rates to valve duties scales the duty linearly to zero below a rate band. The proportional swing rate of a lateral correction of 0.1 m at 7 m radius (about 0.02 rad/s) lies below a band of 0.04 rad/s and would be commanded below the valve opening; the swing's band is therefore 0.01 rad/s (an ablation in Section 10.3 uses 0.04 rad/s). Second, once the teeth are in the soil, the swing hardly corrects a lateral error of the bucket. Every pass, the first one after the return from the dump place included, therefore starts with the reach sub-phase, which brings the lip to the pass start and onto the strip line (within 0.05 m, the error at which the proportional swing rate falls below the band) before the penetration; the return swing ends anywhere over the strip and rebounds when it stops. A reach that has not started the penetration after 60 s without a protective stop or speed limit ends the pass, which counts as a pass without decrease: the start of a pass can lie where the bucket cannot reach it without entering a ground row (a finishing pass against the far batter), and the reach has no time limit of its own. Third, the obstacle rows of the safety layers keep the bucket 0.5 m from every static obstacle (more in the perceived world, whose distance field errs toward the machine), so an object near the far corners of a pit keeps the bucket away from them, and the far half of the side walls would stay above the grade; the pit cones of the site stand 2 m behind the marking (Section 3.3).

The depth that the teeth reach at the far wall is set by the heel of the bucket. In the cutting attitude of the drag (bucket pitch −95°), the rear shell of the bucket lies 0.7 m beyond and 0.75 m above the teeth, over the uncut ground behind the far batter, and the ground row of the bucket admits the teeth only 0.13 m above the grade until they are 0.4 m inside the far end of the flat admissible floor; the far boundary row of the pit lies 0.3 m inside it (Fig. 4). From the barrier geometry, the most curled attitude at which the ground row still admits the teeth at the grade is −80° anywhere on the flat floor, −95° from 0.4 m and −110° from 0.7 m inside its far end. On the grade layer, the drag therefore keeps the bucket open near the far wall and curls it as the teeth move away: its pitch is the more open of the drag's curl schedule (with the pass progress and the fill) and a limit of −80° within 0.1 m of the far end, lowered by 45° per metre beyond. Within 0.45 m of the far end, a lip more than 0.06 m above its target is pulled in with a quarter of the radial intent, so that the teeth reach the floor before the drag leaves the far wall. A drag ends as stalled when its lip has moved less than 2 mm per control step for 1 s (soil resistance, a barrier, the pit wall); in this zone, where the lip is pulled in slowly, the threshold is 0.5 mm.

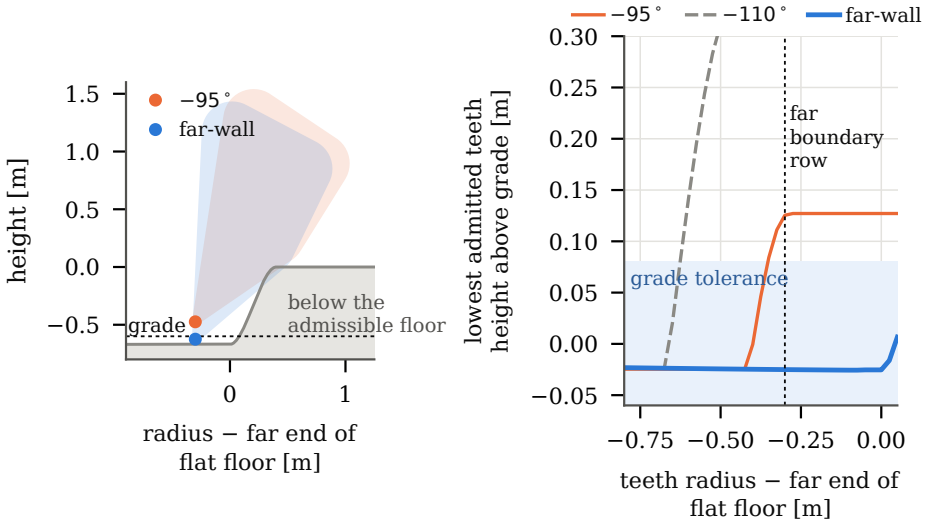

Figure 4: Bucket attitude at the far wall of the grade layer, from the ground row of the bucket and the admissible floor of the last layer. Left: silhouette of the bucket as the row sees it (teeth and shell spheres) with the teeth at the far boundary row of the pit, curled to −95°, where the heel over the far batter holds the teeth above the grade, and in the far-wall attitude, at the grade. Right: lowest teeth height that the row admits, above the design grade, against the radius of the teeth relative to the far end of the flat floor.

The direction of every loaded swing is chosen when the carry starts: the shorter direction whose swept sector (bucket at the carry radius and height, arm, counterweight) keeps the obstacle standoff plus 0.2 m, unless that sector passes within 6 m of a tracked worker or a reachable set and the other direction keeps at least 1 m more, since a fast swing toward a person is slowed or stopped by the separation monitor. The loaded swing starts as soon as the bucket is 1 m above the ground at the pit and waits for the dump clearance only within 1.2 rad of the dump place. A loaded swing that has not progressed by 0.05 rad in 8 s, more than 0.3 rad from the dump place, outside a protective stop and not waiting for the truck to park, is blocked. Without a speed limit of the separation monitor, it continues in the other direction. A speed limit slows a free swing but does not stop it, so a swing that does not progress under a limit is held by something else, typically a spoil heap between the pit and the dump place that is higher than the clearance, which is planned above the heap at the dump point: the bucket is then lifted by 0.3 m with the swing held, up to 1.2 m above the clearance, on the way to the dump place (within 0.3 rad of it the planned clearance holds, at which the dump point lies within the reach), and after the last lift the direction alternates as well. On the way back, the empty bucket only has to clear the dump place itself (load top or heap plus 0.4 m; when the lip has not risen by 2 cm for 3 s, as with a high truck load at the carry radius, the return swing is released); the obstacle rows of both safety layers keep the bucket clear of the truck and the obstacles in between.

7.4. Dump planning and truck exchange

The bucket enters its dump window with the bucket curled, and the largest lip radius of the curled bucket at a given height follows from the front kinematics over the joint ranges (7.0 m up to 2 m above the ground and 6.6 m at 4 m). The dump point is chosen within this reach, 0.3 m inside it. On the marked spoil area, the candidates lie on a 0.2 m grid 0.8 m inside its edges; at the start of every loaded swing the reachable candidate with the lowest heap is selected, so that the heaps spread along the area instead of rising until the dump window leaves the reach. A loaded swing that stalls short of the dump window for 4 s moves the dump point 0.3 m toward the machine (at most 1.2 m). On the truck, the bed centre is used when it is within the reach at the carry height, otherwise the point of the bed's centre line within 1.5 m of the bed centre that is nearest to the machine. A truck that cannot take the soil of the next pit (a budget of 4.2 m^3 of loose soil, which covers the batter rings, the pass margins and spill at the swell factor 1.25) leaves the site and is replaced by an empty truck, which approaches the loading position of the current pit; the machine digs while the truck is not parked and holds a loaded bucket until it is.

8. Learned Residual Skills

Every skill of the system has a structured prior: pure pursuit with final heading alignment for travel, the scripted dig–dump cycle for excavation, and the deadzone-aware inverse valve model with rate feedback for actuation. For travel and excavation, a learned *residual* policy adds a bounded correction to the prior instead of learning the behaviour from zero [36, 37]. Whether learning adds value on top of carefully designed priors and safety filters is an empirical question; this section describes how the residuals are trained and screened, and Section 10.7 reports whether they improve on their priors. The evaluated system of Sections 10.3–10.6 uses the priors.

8.1. Policies

Each policy is a multilayer perceptron with Gaussian outputs (two hidden layers of 256 units, 128 for navigation), trained with proximal policy optimization (PPO) [31] and generalized advantage estimation [32] on the reduced-order twin, with a critic that also receives privileged simulator quantities. Training starts at the prior's behaviour: the output layer of the excavation residual is initialized near zero, and the navigation residual is first fitted to a zero correction on states visited by the prior. The *navigation residual* (scale 0.5) issues travel-speed and yaw-rate corrections toward a stance goal; it observes the goal in the base frame, the heading error, the measured speed and yaw rate, 16 range features to the static obstacles (the analytic obstacle shapes of the reduced-order twin), the nearest worker's relative position and velocity, its previous action and the pure-pursuit intent, and it is rewarded for progress and for arrival (within 0.3 m and 5° of the goal at a speed below 0.1 m/s), with penalties for time, action changes, safety-layer interventions, violations and the residual magnitude. The *excavation residual* (scale 0.5) corrects the task-space intents of the scripted cycle; its reward combines progress toward the phase targets and the fill, transition bonuses (the dump bonus grows with the delivered volume) and penalties for time, spill, interventions, tool forces, intent changes, violations and the residual magnitude. Training environments randomize the plant—for navigation, dead times and lags ±30 %, valve gains ±15 %, friction ±30 % and deadzones ±3 duty points; for excavation, the wider ranges of the stress domain used by the gate below (dead times and lags ±45 %, valve gains ±20 %, friction ±45 %, deadzones ±4.5 duty points); for both, also the bulk modulus, pump margin and delivery, engine speed, sensor scale and ground support—, the soil (±25 %), the dig and dump places, obstacles and walking workers. Both safety layers run during training. Every training run has a fixed budget: 240 PPO iterations of 64 environments × 64 control steps for navigation (983,040 interactions) and 200 iterations of 128 × 64 steps for excavation (1,638,400 interactions).

8.2. Checkpoint gating

A checkpoint gate screens the policy updates before they are accepted (Algorithm 1): every 20th (navigation) or 10th (excavation) iteration, the candidate is verified on the same cohort of seeded stress scenarios (32 trials of widened plant and terrain ranges, its own seed), each of the 32 parallel environments contributing exactly one trial, with unfinished trials in the denominator. An update is accepted only if (i) the violation count in the self-collision, obstacle, worker and footprint categories does not increase, (ii)

Algorithm 1 Fixed-cohort checkpoint-gated PPO

Require: training environments $\mathcal{E}$; verification seeds S_v; period M; initial policy and optimizer state X_0

$X \leftarrow X_0$, $X^\star \leftarrow X_0$, $C^\star \leftarrow \textsc{Verify}(X_0, \textsc{FreshEnvironments}(S_v))$
for $k = 1, \dots, K$ **do**
  collect shielded rollouts in $\mathcal{E}$ and update X by PPO
  if $k \bmod M = 0$ or $k = K$ **then**
    $C \leftarrow \textsc{Verify}(X, \textsc{FreshEnvironments}(S_v))$
    if $\textsc{Gate}(C, C^\star)$ **then** $X^\star \leftarrow X$, $C^\star \leftarrow C$
    else restore $X^\star$; reduce the learning rate after two consecutive rejections
    end if
  end if
end for
evaluate $X^\star$ once on independently seeded held-out cohorts

the count of reference-layer infeasibility does not increase (navigation) or at most doubles plus ten steps (excavation, where soil contact makes it frequent), (iii) the joint-plus-ground violation count increases by no more than the larger of two and 10 % of the incumbent's, and (iv) the task performance on the cohort does not fall below that of the last accepted policy: the completion rate for navigation (a trial is complete when the environment reports arrival at the stance; ties broken by the shorter mean trial time) and the delivered volume for excavation. A rejected candidate restores the accepted actor, critic, observation normalizers, optimizer moments and optimizer random state, and two consecutive rejections reduce the learning rate. The task performance of the accepted policy on the verification cohort can therefore not fall below that of the initial policy, which reproduces the prior up to its small initial residual, and its joint and ground violations may grow only within the tolerance of (iii); independent held-out cohorts, never used for selection, measure whether it is better.

9. Experimental Protocol

The control-layer benchmarks (Section 10.2) and the learning studies (Section 10.7) use the reduced-order twin; all other experiments run on the full twin of Section 3 with the control stack of Sections 4–7, the scripted cycle, pure pursuit and the inverse valve model. The design parameters were fixed before the experiments reported here (configuration file in the code archive) and are the same in all of them. They were developed on the twin with diagnostic runs of single pits (P3 and P8 with seeds 101 and 102, P2 with seeds 401 and 402, P3 with seed 402) and of the missions of seeds 401–420, in parts and, in six earlier runs described below, as a whole; results for these pits and seeds are therefore in-sample, whereas the missions of seeds 421–425 were not used in the development. Two perception conditions are used: *perceived*, in which every layer receives only the outputs of the perception stack of Section 5, and *privileged*, in which the exact soil surface, obstacle geometry and actor states replace them (used to isolate the task layer from perception, and as the oracle condition of the worker study). The evaluation itself always uses the true state. In the missions and the single pits (with their ablations and sensitivity runs), a physics audit evaluates the barrier rows with the exact obstacle geometry, the true positions of the actors and the admissible floor of the true work-order rectangles, and counts the samples in which a safety-critical row (self-collision, obstacle, worker, footprint) lies below zero; the worker study evaluates the clearance to the worker directly (below).

Model checks (Section 10.1) compare the full twin with the manufacturer's specification sheet of the DX300LC-5 [50]: working ranges from the front kinematics over the joint ranges; bucket and arm digging forces according to ISO 6015 [51] (cylinder force at the relief pressure converted to the tooth tip; largest over the stroke); swing and travel speeds at full duty from rest (the swing empty and with a $1\,\mathrm{m}^3$ payload); and the bucket capacity. The soil model is examined by the sensitivity of pit completion to its parameters (Section 10.4).

Single pits (Section 10.3) excavate one pit from its stance with the site actors present, in the privileged world (the perceived pits are those of the ten-pit missions): P3 (truck) and P8 (spoil), two scenario seeds, a time limit of 900 s per pit. Design ablations remove one element of the task layer at a time (P3 and P8,

seed 101): the stable-wall layering (one layer cut to grade), the pass margins on the far and near side, the far-wall attitude of the grade-layer drags (the bucket curled to the cutting attitude from the pass start), the fine rate band of the swing (0.04 rad/s instead of 0.01 rad/s) and the dump planning (a fixed dump point on the centre line of the spoil area opposite the pit, or the bed centre of the truck). The single pits and the ablations use pits and seeds that were also used in the development of the design; they show mechanisms and their effect on one scenario, not completion rates.

Sensitivity (Section 10.4) repeats the privileged single pit P8 (seed 101) with one soil parameter changed at a time: the bank cohesion halved and doubled (2.45 kPa and 9.8 kPa), the friction angle 32° and 42°, the steepest standing wall of bank soil 60° and 80°, and the bank density 1,400 kg/m^3 and 1,800 kg/m^3; and the pump flow of the plant reduced by 10 % (P3 and P8; the controller and the safety layers keep the nominal model). For the 60° wall, the run is repeated with layers and batter rings of 0.15 m, the step that this wall admits.

Ten-pit missions (Section 10.5) execute the complete work order P1–P10 with perception in the loop, from the start position at the west end of the site: survey of the site, pit detection, travel, excavation to grade, truck loading and exchange, spoil dumping. Ten scenario seeds vary the pauses of the site workers, the noise and dropouts of the LiDARs, the dropouts of the near-field scanners and the random spill of the soil model: the development seeds 401–405 and the held-out seeds 421–425, which were chosen before the experiments reported here and used in no other run (a second campaign with the same frozen design and sources). A pit ends when the planner judges it complete from its estimate, when the remaining material of no strip is reduced any more, or after 1,800 s of excavation (when the bucket is next out of the soil and empty, at the latest 60 s later); the mission ends when the work order is processed or after 11,000 s.

Six earlier runs of the missions shaped the final design and site, three of them on seeds 401–405. In the first, with the pit limit of the single pits (900 s), P10 of three missions reached the limit unfinished, and the pit limit was raised to 1,800 s, because people near the machine can hold it for long periods. The second run showed two defects of the simulated site at its east end. First, the site workers, stepping back from the approaching machine as the site rule demands, reached the end of the walkable site, where the model of the site workers had no way out of the zone, and stood there during 33–95 % of the excavation time of P8 and, in every mission, throughout the excavation of P9 and P10, 11.8 m to 14.2 m from the swing axis (median), inside the zone. Second, the surrounding terrain rises to the ground level inside the border of the soil heightfield, across P10: once the pit was dug, the simulated LiDARs returned that surface, 0.26 m to 0.57 m above the grade, instead of the pit floor, so that, in the three missions that reached the grade layer of P10, the estimate of the strip along the pit's west edge kept returning to about 0.26 m above the grade and the planner retried that strip until the limit. We added the exit rule of Section 3.3 and the crossing of W2 behind the machine from P6 on to the model of the site workers, and let a LiDAR ray whose first hit is the surrounding terrain more than 5 cm above the soil be cast again without it (on the undug site, the scans are unchanged). The third run, with these corrections, exposed two liveness gaps of the design at P8, the pit in front of the pipe stack. In one mission, a near-field detection on the end face of the pipe stack, which from the stance of P8 only a scanner sees, held the machine in a protective stop until the pit's time limit; in another, the bucket came to rest between the ground row and the standoff of the pipe stack at the start of a finishing pass, where the actuator layer admitted no motion. The column rule of the near-field static test (Section 5.3), the relaxed condition of the actuator layer at rest (Section 4.4) and the reach watch (Section 7.3) close these gaps. The fourth run, with these rules, on seed 401 and on seeds 406–410, which had been held out until then, exposed one more gap, at P9 of seed 407: a loaded swing toward the spoil area came to rest at the standoff of a spoil heap that was higher than the carry clearance and stayed there, advancing by 0.05 rad in about 500 s, until the run was stopped. A false person track 12 m east of the machine, where no worker was, limited the speed, which suspended the recovery of a blocked swing, and the relaxed condition of the actuator layer gave commands inside the valve deadzone, which did not move the machine. The lift of a loaded swing at rest under a speed limit (Section 7.3) closes this gap, and seeds 406–410 thereby became development seeds. The fifth run, with the lift, on seeds 411–415, which had been held out until then, was stopped after P3 of seed 413 had ended by the grade check with one cell 0.13 m above the grade. A replay from the state saved at the start of the pit traced the miss to a load of 0.045 m^3 that a finishing pass had left in the bucket: it fell out when the reach of the next pass opened the bucket, and the part beside the cut of that pass stayed in the near rows, which the LiDAR does not see;

the same mechanism had left P4 of seed 401 8.3 cm above the grade in the fourth run. The grade-layer carry (Section 7.2) closes this gap: in replays from the saved states, both pits reached the grade with it (largest deviations 3.6 cm and 6.9 cm). Seeds 411–415 thereby became development seeds as well. The sixth run, with the carry, on seeds 416–420, which had been held out until then, exposed a gap in the tracker. In the mission of seed 417, a person-sized segment of returns from the front face of the pipe stack behind P8 became a track whose estimate moved more than 0.6 m from where it was first seen while most of its points were not yet long-observed static surface, which the tracker takes for a walking person (Section 5). The status was never revoked, although the track's segment lay on long-observed static surface from then on: the separation monitor held the machine in a protective stop through the time limit of P8 and then in the fold for P9, until the run was stopped. The revocation of this status on static structure (Section 5) closes this gap: in a replay from a state saved during P8, the track was revoked 10 s after the start of the replay, the machine moved at once, and the pit ended by the grade check at grade (largest deviation 4.2 cm). In the same mission, P3 had ended by the grade check with one cell 9.3 cm above the grade. A replay from a state saved shortly after the start of the pit, which reproduced it, traced the miss to the displacement rule of the soil model (Section 3.1): in the first quarter second of the curl at the near end of two late passes, the rule moved 5.4 l of soil that the bucket shell overlapped beyond the near wall as far as 1.8 m along the shell's motion, across the bucket's footprint, onto floor near the far wall that an earlier pass had cut to the grade, and the LiDAR estimate of that cell rose to 6.6 cm only. This is a property of the twin's soil model, which the design does not address. The missions of seeds 416 and 418 completed the work order, with 19 of their 20 pits at grade; one cell of P10 of seed 416 ended 0.2 cm outside the tolerance. Seeds 416–420 thereby became development seeds as well, and the held-out missions use the new seeds 421–425. The results reported are those of the final design and site: the worker study and the missions were run with them, and the earlier runs of the missions are summarized in Section 10.5. The single pits, the ablations and the sensitivity runs use exact perception, in which the mission has no tracker; they were run before the revocation was added, with the final design otherwise, and a rerun of one of them (P8, seed 101) with the final sources reproduced it exactly (its summary, events and final soil surface). The LiDAR correction does not affect any other experiment, since none of them digs at the east end of the site. The controlled stops of the stopping-bound identification (Section 6.1) were not repeated: they are commanded at random times with the planner held, the relaxed condition of the actuator layer does not act during a stop, and the site workers and the other rules affect only the states at which the stops were sampled.

Worker protection (Section 10.6) uses the stopping-bound identification of Section 6.1 and 50 scenarios in which one worker follows a prescribed, time-parameterized path that ignores the site rule and the machine, while the machine excavates P3 with truck loading (scenario seed 301). Every path starts 18 m from the swing axis, outside the 16 m region of condition (E2); the worker walks to the start point of its family, stands there until a random time of up to 45 s and then (A) walks into the path of the loaded bucket between the pit and the truck, stands there and walks back; (B) crosses between the machine and the truck, under the loaded swing; (C) walks to the far edge of the pit being dug, stands and walks back; or (D) crosses between the travel line and the pit, under the boom and through the tail swing; each at 0.8 m/s to 1.4 m/s; (E) is (A) at running speed (2.0 m/s to 2.6 m/s), outside the envelope. Ten scenarios are drawn per family. Five conditions use the same scenarios, plant and seeds: separation monitor with exact worker states, with perception, and with perception and a 1 s blackout of all object returns around the encounter; both safety layers without the monitor (exact states; workers with their velocity in the rows, which restore their standoff); and neither (exposure of the scenarios). Clearance is evaluated at the worker's exact position every 5 ms: horizontally between the worker disk (radius 0.35 m) and the collision circles of the machine, and in three dimensions between the worker body (a capsule from the ground to 1.8 m) and the collision spheres of the front and upper structure and the track boxes. Each sample is classified as *machine-induced* when the machine's own motion since the previous sample reduced the clearance to the worker's current position at more than 0.05 m/s, and as *worker-induced* otherwise; the smallest clearance is reported over all samples and over the machine-induced ones. In addition, every reduction of the clearance by the machine's motion while the worker is inside the standoff, however slow, is summed over the trial (the machine's share of the closing).

Learned residuals (Section 10.7) are trained with five seeds each and evaluated once, together with their priors, on held-out cohorts of 64 trials with nominal and with widened (stress) plant and terrain ranges.

Differences are reported with 95 % confidence intervals over the five training seeds (Student's t); they do not include the sampling variation of the fixed cohorts.

10. Results

10.1. Model checks against the manufacturer's data

Table 2 compares the full twin with the specification sheet of the machine [50]. The working ranges follow from the calibrated front geometry and the joint ranges and agree within 0.4 %. The digging forces were not used in the calibration: they combine the cylinder areas, the relief pressures of the specification and the moment arms of the calibrated linkage, and they are 2.4 % to 3.9 % below the rated values, i.e., the linkage transmits the cylinder forces as the manufacturer states. The swing and travel drives were sized from the same sheet, but at full duty the twin swings at 7.58 instead of 9.7 rpm (22 % slower; 90° from rest in 2.65 s, 180° in 4.65 s) and travels at 2.60 instead of 3.0 km/h in the low range (13 % slower); the valve curves of these drives were assumed, not measured, and the twin is used as specified, without compensating these deviations. Swing and travel times are therefore longer than on the machine, while the stopping bounds of Section 6.1 are those of the twin's slower swing: on the machine they would have to be identified again with the same procedure. The bucket model is larger than the machine's bucket: it holds 1.22 m^3 struck and, with a heap of 1:1 slopes over its strike plane (1.81 m × 1.38 m; the heaped rating of ISO 7451 [52]), 1.86 m^3, 46 % above the rated heaped capacity of 1.27 m^3; its cavity is modelled as the convex hull of the bucket shell, which overestimates the cavity of the real bucket. The loads of a dig cycle can therefore be larger than on the machine, so that the number of passes per pit is not conservative, and the pass sizing of Section 7.2 refers to this bucket.

No digging data were recorded on the machine, so the soil model and the digging loads are not validated against measurements. Section 10.4 therefore measures how the completion of a pit depends on the soil parameters, on the wall stability and on the pump capability; it changes one parameter at a time on one pit and does not show that the results hold for every soil within these ranges.

Table 2
Full twin versus the specification sheet of the DX300LC-5 [50]: digging forces according to ISO 6015 in normal and power mode, heaped bucket capacity according to the SAE rating of the specification (struck capacity of the bucket model 1.22 m^3). Deviation in percent of the specified value; the last column names what the twin's value depends on.

Quantity	Specification	Twin	Dev. [%]	Depends on
Largest reach at ground level [m]	10.530	10.536	+0.1	kinematics, joint ranges
Largest digging depth [m]	7.305	7.300	−0.1	kinematics, joint ranges
Largest digging height [m]	10.325	10.362	+0.4	kinematics, joint ranges
Bucket digging force, normal / power [kN]	185 / 196	180.5 / 190.8	−2.4 / −2.7	cylinder, relief pressure, linkage
Arm digging force, normal / power [kN]	129 / 136	123.9 / 131.0	−3.9 / −3.7	cylinder, relief pressure, linkage
Swing speed at full duty [rpm]	9.7	7.6	−21.9	swing motor, pumps, valve map
Travel speed, low range [km/h]	3.0	2.6	−13.4	travel motors, pumps, valve map
Bucket capacity, heaped [m^3]	1.27 (SAE)	1.86	+46.3	bucket geometry

10.2. Control layers on the reduced-order twin

Table 3 evaluates the safety layers on 128 randomized machines of the reduced-order twin, driven for 20 s each by random smooth intents through the planner (51,200 control steps per configuration; joint, self-collision and ground rows; random payloads and external torques). Without filters the machine violates the barriers in 7,132 steps, up to 497 cm below the admissible floor. The reference layer alone reduces this to 1,980: a filtered reference is not enough when the valves deliver it late, partially or not at all. Actuator-layer variants that credit corrective commands in full (linear condition, Eq. (5)) remain at several hundred violations whatever dead-time prediction they use (none; the model's first-order relaxation toward the in-flight command over the dead time; continued motion with the worst-case acceleration toward the barrier; or the interval prediction of Proposition 2), consistent with corrective motion being withheld by the deadzone, shared pump flow or a stalled function. The braking credit removes most of them (594 with the

linear condition and the interval prediction, 31 with the credit), the interval prediction reduces the violations of the credited condition further (88 without dead-time prediction), and the rate-consistent reference layer improves on a velocity-level reference (67) and on no reference filtering (72); without reference filtering the actuator layer intervenes about three times as much (mean ι_a 0.057 versus 0.019). The remaining 31 steps (0.06 % of the samples) are 8 joint-limit steps at a cushioned end stop (inside the 1° joint margin) and 23 ground steps, the deepest 15.0 cm below the floor.

Table 3
Ablation of the safety layers (128 machines × 20 s, model-based actuator controller, random smooth intents through the planner, payloads and external torques). Violations are counted in control steps; "Crit." sums the safety-critical categories present in this benchmark (self-collision); "Depth" is the deepest ground penetration; ι_a the mean actuator-layer intervention (sum of the absolute changes of the commanded coordinate rates per control step); tracking error in units of $u_{\max}$. Bold: the configuration of the system.

Reference layer	Actuator layer	Viol.	Joint	Crit.	Ground	Depth [cm]	ι_a	Track err.
none	none	7,132	2,016	795	5,794	496.6	0.000	0.232
rate-consistent	none	1,980	35	75	1,872	42.2	0.000	0.254
none	interval + credit	72	7	0	65	13.5	0.057	0.240
velocity-level	interval + credit	67	2	0	65	15.0	0.018	0.225
rate-consistent	no dead time, linear	717	5	5	709	36.5	0.006	0.242
rate-consistent	model prediction, linear	960	5	20	935	33.3	0.007	0.251
rate-consistent	worst case, linear	426	2	0	424	32.1	0.011	0.244
rate-consistent	interval, linear	594	1	0	593	35.2	0.011	0.242
rate-consistent	credit, no dead time	88	8	0	80	16.7	0.013	0.244
rate-consistent	**interval + credit**	**31**	8	0	23	15.0	0.019	0.243

Table 4 compares the synchronized planner with plan-then-track (the plan integrates the intents in its own time and a feedback term tracks it) and replanning (the plan restarts from the measured state at every control step), under random intents (64 machines, 30 s) and in the scripted excavation cycle with obstacles and workers (32 machines, 120 s). With plan-then-track the plan error grows whenever the machine is held back by a barrier, a stall or the pump limit (from 0.64 in the first to 3.32 in the last 5 s of the episodes, maximum 13.6); the catch-up produces large corrective references, and in the excavation cycle the violations rise to 286 steps, 8 of them safety critical with penetrations of up to 8.9 cm (synchronized: 49 and 2, at most 0.9 cm), at 0.69 instead of 0.73 m^3/min. The synchronized planner keeps the error bounded, consistent with Proposition 1 (95th percentile 1.21, growth 0.15). Replanning keeps the plan error at the motion of one control step (mean 0.08) but discards the plan state: under open-loop intents it produced 147 violation steps versus 85, because nothing slows the reference when the machine falls behind. In the closed-loop excavation cycle, where the scripted controller corrects the position itself, replanning produced fewer violation steps than the synchronized planner (34 versus 49, 0 versus 2 safety critical) at a slightly lower rate (0.70 versus 0.73 m^3/min). The advantage of the synchronized planner is therefore specific to intents that act on the plan open loop, as those of the learned skills, and to the comparison with plan-then-track; for a closed-loop scripted controller, replanning was at least as safe in this benchmark. The system uses the synchronized planner for all skills, so that the scripted priors and the learned residuals of Section 8 run on the same stack; for the evaluated system, which uses the scripted priors, replanning would have been an equally safe choice in this benchmark.

10.3. Single pits and design ablations

Table 5 lists the single pits of the privileged world, and Fig. 5 shows their final surfaces. With the full design, 4 of 4 pits reached the grade at every point and were ended by the planner, after 388–395 s of excavation (largest deviation 5.1 cm). With one element removed, the pits of seed 101 (392 s for P3 and 392 s for P8 with the design) also reached the grade at every point; what changed was the time they took and their margin to the tolerance.

Without the pass margins, P3 took 691 s and P8 545 s (76 % and 39 % longer). With the bucket curled at the far wall, P3 took 589 s and ended 7.9 cm above the grade, at the edge of the tolerance (5.1 cm with

Table 4
Planner–controller synchronization under random intents and in the scripted excavation cycle: weighted plan error $\bar{e}$ (mean), e_{95} (95th percentile) and $\hat{e}$ (maximum), Δe the mean error in the last minus the first 5 s, violation steps (Viol.), safety-critical violation steps (Crit.) and excavation rate.

	Random intents (64×30 s)					Excavation (32×120 s)		
Planner	$\bar{e}$	e_{95}	$\hat{e}$	Δe	Viol.	m^3/min	Viol.	Crit.
synchronized	0.62	1.21	2.2	+0.15	85	0.73	49	2
plan-then-track	2.10	6.91	13.6	+2.68	123	0.69	286	8
replanning	0.08	0.16	0.3	+0.00	147	0.70	34	0

Table 5
Single pits in the privileged world (design, two scenario seeds) and design ablations (seed 101): completion to grade at every point, area within the tolerance, largest deviation from the design grade, unexcavated and overexcavated volume and the excavation time until the planner ended the pit (time limit 900 s; –: not ended within the limit).

Case	Pit	Seed	Complete	Area [%]	Max. dev. [cm]	Unexc. [m^3]	Overexc. [m^3]	Time [s]
design	P3	101	yes	100.0	5.1	0.004	0.040	392
	P3	102	yes	100.0	4.2	0.006	0.036	395
	P8	101	yes	100.0	4.6	0.001	0.045	392
	P8	102	yes	100.0	4.8	0.002	0.062	388
no stable-wall layers	P3	101	yes	100.0	3.9	0.002	0.042	483
	P8	101	yes	100.0	4.7	0.003	0.046	519
no pass margins	P3	101	yes	100.0	5.8	0.006	0.040	691
	P8	101	yes	100.0	5.1	0.007	0.044	545
curled at the far wall	P3	101	yes	100.0	7.9	0.025	0.027	589
	P8	101	yes	100.0	5.2	0.002	0.058	430
coarse swing band	P3	101	yes	100.0	5.6	0.002	0.046	402
	P8	101	yes	100.0	6.3	0.001	0.067	385
fixed dump point	P3	101	yes	100.0	5.1	0.004	0.040	392
	P8	101	yes	100.0	5.1	0.001	0.066	407

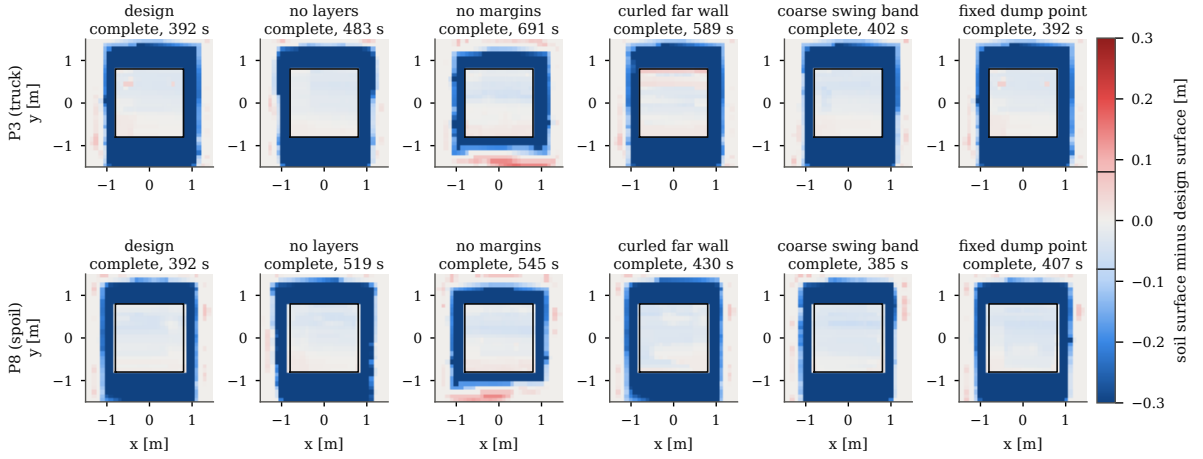

Figure 5: Final soil surface of the single pits P3 and P8 (seed 101, privileged world) with the full design and with one element of the task layer removed, relative to the design grade inside the rectangle (outlined) and to the original ground outside it, where the stepped walls of the pit appear as excavation; points outside the grade tolerance (± 0.08 m, the lines on the colour bar) are marked, and the colour scale saturates at ± 0.3 m.

the design), and P8 took 430 s. Without the stable-wall layering, P3 took 483 s and P8 519 s (23 % and 32 % longer). The coarse swing band changed little in these pits (402 and 385 s; largest deviations 5.6 and 6.3 cm). With a fixed dump point, P3 repeated the run of the design in every pit metric, since the bed centre of the parked truck lay within the reach of the curled bucket and the planner chose it as well, and P8 took 407 s; a single pit does not test the spreading of the spoil of several pits over the spoil area, for which the planning was designed.

10.4. Sensitivity to the soil and the pump

Table 6 repeats the privileged single pit P8 with one soil parameter changed at a time and both pits with 10 % less pump flow. With every soil variant, the pit was complete to grade within the time limit, after 386–533 s of excavation (392 s with the nominal soil): with half and double the bank cohesion (largest deviation 4.8 and 5.2 cm), with friction angles of 32° and 42° (6.6 and 4.0 cm), with bank densities of 1,400 kg/m^3 and 1,800 kg/m^3 (7.4 and 5.0 cm) and with a steepest standing wall of 80° and 60° (6.9 and 4.5 cm). With the 60° wall, the 0.2 m layers of the design exceed the step at which the wall fails; the pit was nevertheless complete after 506 s, and with the 0.15 m layers that this wall admits after 510 s. With 10 % less pump flow, both pits were complete, after 495 s of excavation at most. Within the ranges examined, the soil parameters changed the excavation time of the pit and its margin to the tolerance, not its completion; one pit and one scenario seed per variant do not show how often a variant would fail.

Table 6
Sensitivity of pit completion (privileged world, seed 101) to the soil parameters and to the pump flow of the plant: completion, area within the tolerance, largest deviation and excavation time.

Parameter	Value	Pit	Complete	Area [%]	Max. dev. [cm]	Time [s]
bank cohesion	2.45 kPa	P8	yes	100.0	4.8	456
	9.8 kPa	P8	yes	100.0	5.2	427
friction angle	32°	P8	yes	100.0	6.6	386
	42°	P8	yes	100.0	4.0	426
steepest bank wall	60°	P8	yes	100.0	4.5	506
	60°, 0.15 m layers	P8	yes	100.0	5.1	510
	80°	P8	yes	100.0	6.9	432
bank density	1,400 kg/m^3	P8	yes	100.0	7.4	432
	1,800 kg/m^3	P8	yes	100.0	5.0	533
pump flow	90 %	P3	yes	100.0	4.9	495
		P8	yes	100.0	6.5	381

10.5. Ten-pit missions

Table 7 summarizes the ten missions with perception in the loop and Fig. 6 their timelines and pits.

99 of the 100 pits reached the design grade at every point of the true soil: 50 of the 50 pits of the development seeds 401–405 and 49 of the 50 pits of the held-out seeds 421–425. Nine of the ten missions completed the whole work order to grade, four of the five held-out missions among them; a mission took 1.50–1.58 h of mission time (survey, travel, excavation, truck exchange and dumping). The pits of a mission share its site, its workers and the state of the machine and are not independent; with the missions as the unit, the 95 % interval of the rate of pits at grade is 97.0–100.0 % (percentile bootstrap over the missions) and that of the rate of complete missions 60–98 % (Wilson).

Every pit was ended by the planner's grade check, after 457–628 s of excavation (limit 1,800 s; 533 s and 10.3 passes on average). The pits at grade kept a margin to the tolerance: their largest deviation from the design grade was 4.1 cm in the median pit and at most 7.2 cm (tolerance 8 cm; Fig. 6, bottom right), the mean deviation over a pit was 1.7 cm, and a pit was overexcavated by 0.038 m^3 on average (design volume 1.536 m^3). The passes followed their strips: at the far end of the passes, where the penetration starts, the bucket tip was 3.1 cm (median; 95th percentile 5.9 cm) from the strip line, and at the near end 3.6 cm (6.6 cm), over 981 passes. Every pit was found by the cameras and every stance was reached; no carry stalled

Table 7
Ten-pit missions with perception in the loop (development seeds 401–405 above the line, held-out seeds 421–425 below it): pits at grade at every point (true soil), completion of the work order, mission time, area within the tolerance over all pits, delivered loose volume, protective stops (duration and count), speed-limited time, smallest clearance between a worker and the machine's collision geometry, and samples in which the physics audit found a safety-critical boundary breached.

Seed	Pits	Order	Time [h]	Area [%]	Deliv. [m^3]	Stop [min]	Stops	Slow [min]	Worker [m]	Audit
401	10/10	yes	1.55	100.00	39.6	3.8	190	7.2	3.64	0
402	10/10	yes	1.53	100.00	39.2	4.0	174	6.7	3.63	0
403	10/10	yes	1.54	100.00	39.2	3.9	186	6.7	3.51	0
404	10/10	yes	1.50	100.00	38.4	4.1	199	7.3	3.50	0
405	10/10	yes	1.58	100.00	40.6	4.7	258	8.2	3.67	0
421	10/10	yes	1.56	100.00	39.6	4.6	203	7.6	4.06	0
422	10/10	yes	1.57	100.00	40.1	5.7	194	6.8	3.52	0
423	10/10	yes	1.54	100.00	38.4	3.7	200	6.9	3.53	0
424	9/10	no	1.54	99.96	38.3	4.7	286	8.0	3.54	0
425	10/10	yes	1.51	100.00	39.3	3.6	183	7.3	3.47	0

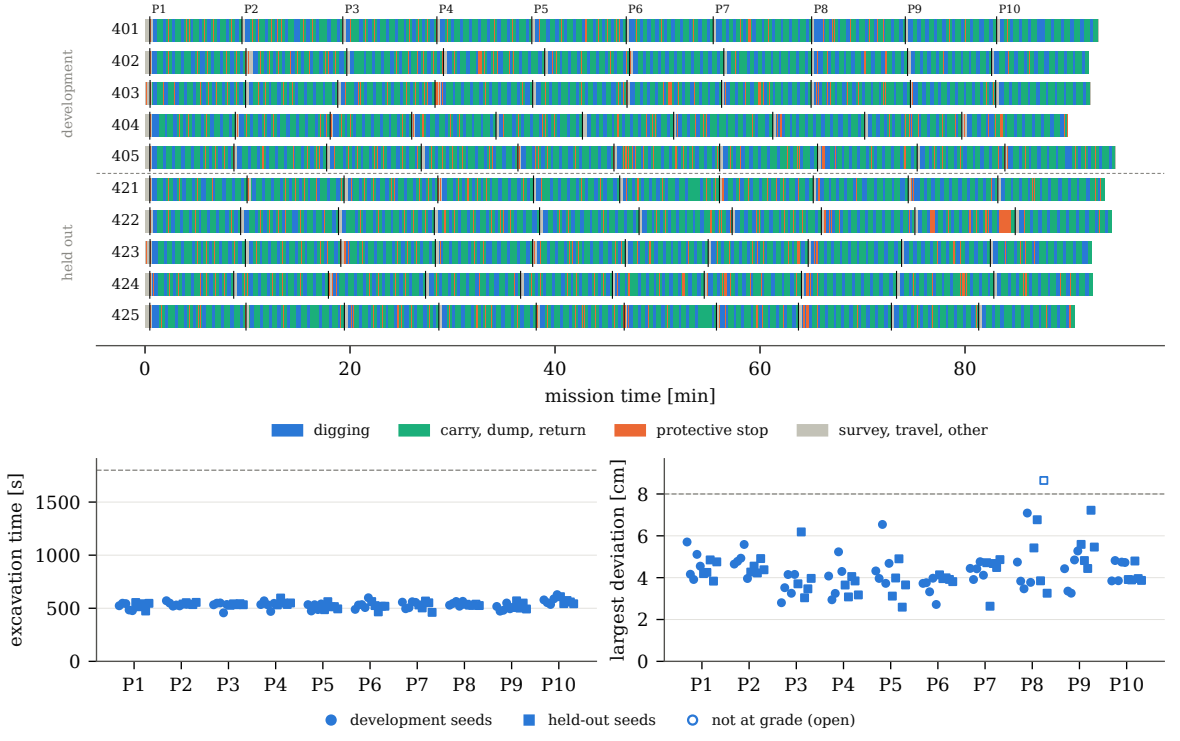

Figure 6: Ten-pit missions. Top: timeline of every mission (digging, carry/dump/return, travel and survey, protective stops), development seeds above and held-out seeds below the dashed line, with the end of every pit marked. Bottom left: excavation time of every pit, with the time limit (dashed). Bottom right: largest deviation of every pit from the design grade on the true soil, with the tolerance (dashed). Circles: development seeds; squares: held-out seeds; filled: complete to grade; open: not complete.

short of its dump window, and no travel was blocked. A pass carried at most 1.09 m^3 of loose soil, and the missions delivered 39.3 m^3 of loose soil each on average.

One pit missed the criterion, P8 of seed 424: one of its 289 cells ended 8.6 cm above the grade, 0.6 cm beyond the tolerance, after the planner's grade check had ended the pit with a largest estimated residual of 5.0 cm. A replay from the state saved shortly after the start of the pit, which reproduced the pit exactly,

traced the miss to the mechanism that had left P3 of seed 417 outside the tolerance in the sixth earlier run (Section 9). Pass 6 had cut the cell, 0.1 m from the side wall and 0.5 m from the far wall, to the grade; in the first half second of the curl at the near end of two later finishing passes, the displacement rule of the soil model moved 5.21 of soil that the bucket shell overlapped beyond the near wall 1.5 m to 1.7 m along the shell's motion, across the bucket's footprint, onto the cell and its neighbours, and at the planner's decision the LiDAR estimate of the cell was about 2 cm above the grade.

The machine was stopped by the separation monitor for 4.6 % and speed-limited for 7.9 % of the mission time. Unlike in the second earlier run (Section 9), the stops did not grow toward the east end of the site: they took 4 % of a pit's excavation time on average in P1–P7, 5 % in P8 and P9 and 3 % in P10, and at most 24 %, in P9 of seed 422, mostly while a site worker paced about 16 m from the swing axis. A site worker was within 16 m of the swing axis during 28 % of a pit's excavation time on average, and while the machine was stopped, the nearest worker stood 16.1 m from the swing axis (median). No site worker came closer than 3.47 m to the machine's collision geometry (Table 7).

The five liveness rules acted as designed. The relaxed condition of the actuator layer moved the machine from rest in 899 control steps (ten missions), the reach watch ended no reach, the column rule of the near-field static test made 42,004 near-field returns static (by design all of them beyond 8 m from the swing axis), 10 loaded swings were lifted at rest under a speed limit (four missions), and the tracker revoked the worker status of 24 person tracks (nine missions). Each of the revoked tracks lay within 0.01 m of the footprint of a static object of the site, the roller (12), the fence (7), the toilet (3) and the pipe stack (2), at least 7.5 m from the machine's stance; the nearest site worker was 2.1 m from a revoked track, and at every revocation the site workers were at least 11.3 m from the machine's swing axis. No mission was held by a safety layer without a hazard until a limit, and the physics audit found no sample in which a safety-critical boundary was breached.

The earlier runs of the missions (Section 9) ended as follows. The first brought 45 of the 50 pits to grade and completed one mission, the second 46 pits and two missions. The third completed the missions of seeds 401 and 402 and stalled at P8 of seeds 403 and 404, after P1–P7 had reached the grade. The fourth completed the missions of seeds 401 and 406, with 19 of their 20 pits at grade (P4 of seed 401 8.3 cm above it), and was stopped at P9 of seed 407 and P10 of seed 408, before the missions of seeds 409 and 410 had started. The fifth was stopped after nine pits of seeds 411, 412 and 413, eight of them at grade, before its other missions had started. The sixth completed the missions of seeds 416 and 418, with 19 of their 20 pits at grade (P10 of seed 416 8.2 cm above it), and was stopped after the machine of seed 417 had been held from P8 on; in that mission, P3 had ended 9.3 cm above the grade (Section 9). The simulation is deterministic: 44 of the 44 pits that the first run had ended before its time limit ended identically in the second, which differed only in the limit, and the missions, which the computing environment interrupted several times, resumed from their saved states without a difference (a replay from a saved state reproduced the recorded pit exactly, also across a resume).

Perception: within 16 m of the machine the tracker reported 72 % of the worker samples with a precision of 93 %; the near-field scanners reported a person in 0.5 % of the samples, and 99 % of these detections were not at a worker; the elevation estimate of the pit area deviated by 5.6 cm (root mean square) from the true soil.

10.6. Worker protection

Stopping bounds. The identification set comprises 60 controlled stops (21 while digging, 11 in loaded carries, 19 in empty returns, the others while approaching or dumping; 20 from circle speeds above 1 m/s, up to 3.5 m/s). Figure 7 shows the travel and the stopping time against the effective speed at the command. A stop from a fast loaded swing travelled up to 2.91 m and took up to 2.25 s until no circle moved by more than 0.05 m; the stopping time is dominated by the decay of the hydraulic oscillation of the upper structure after the swing has stopped, and it grows with the speed up to the cap $T_{\max}$, which $T(v_e)$ reaches at 0.91 m/s. The slope k_T is set by a slow stop, commanded at 0.47 m/s, that took 1.15 s to settle. The motion reference was faster than the circles in 61 of the 100 stops, and the circles were speeding up beyond a_{acc} in 25: the reference leads the hydraulic motion, and a stop commanded at the onset of a motion still lets the machine accelerate. The largest effective speed, 4.1 m/s (set by the reference term), belongs to a stop during a carry: the bucket circles moved at 3.45 m/s and the reference at 4.07 m/s when the stop was commanded, and the

circles still reached 3.68 m/s and travelled 2.91 m. One stop commanded at 0.22 m/s, with the reference at 1.50 m/s, reached 1.10 m/s after the command.

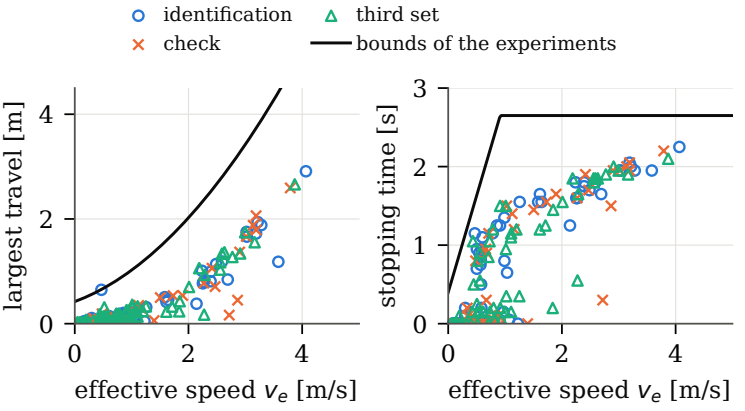

Figure 7: Controlled stops of the full twin: largest travel of a collision circle (left) and stopping time (right) against the effective speed at the stop command. Circles: identification set; crosses: check set; triangles: third set, measured after the bounds of the experiments (lines) were fixed.

The bounds identified on this set ($T_a = 0.40$ s, $k_T = 2.46\,\mathrm{s^2/m}$, $T_{\max} = 2.65$ s, $d_0 = 0.30$ m) held for all 40 stops of the check set, which were measured with other scenario seeds of the two pits: 40 within the time bound, 40 within the travel bound and 40 with every circle within its share. The same rule applied to all 100 stops gives the identified values again; these are the bounds of the experiments. They held for all 60 stops of a third set, measured after they were fixed in three further runs of the two pits with other scenario seeds and the configuration of the experiments: every stop was within the time and the travel bound and had every circle within its share (largest travel 2.66 m, longest stopping time 2.10 s).

No stop was commanded with both the circles and the reference at rest, so T_a is the time margin alone. With the speed at the command instead of the effective speed, a stop commanded at the onset of a motion is evaluated at the low speed of the circles but lasts as long as one from the speed of the reference, and the same rule gives $k_T = 10.57\,\mathrm{s^2/m}$ (with the acceleration term but without the reference, $4.88\,\mathrm{s^2/m}$): the time bound then reaches its cap at 0.21 m/s instead of 0.91 m/s, and a person needs up to 3.0 m more separation from a slow circle (at 0.2 m/s, straight approach, Eq. (10); 0.2 m more at 1 m/s). Of the 100 stops, 58 were commanded below 1 m/s. Both $d(v_e)$ and $T(v_e)$ concern the stopping performance of the twin; on the machine, whose swing is faster at full duty (Section 10.1), they would have to be measured again.

Worker study. Table 8 lists the clearances of the 50 prescribed encounters under the five conditions, and Fig. 8 shows the smallest clearance of every encounter. With the separation monitor, the machine's own motion brought a worker inside the standoff in none of 120 runs: the 40 encounters inside the envelope under each of the three monitor conditions (exact states, perception, perception with blackout; 95 % upper bound of the rate per condition 7.5 % by the rule of three, since the conditions repeat the same encounters), and no three-dimensional contact was machine-induced. The smallest machine-induced clearance was 4.39 m with exact states, 4.16 m with perception and 4.16 m with the blackout, and while a worker was inside the standoff, the machine moved toward the worker by at most 7 mm in a trial.

The smallest clearances overall were worker-induced: the workers of the study ignore the machine and walk under its raised front and through its tail swing. With exact states, the paths of 3 encounters (families B and D) took the worker into the horizontal collision geometry of the machine after the monitor had stopped it (smallest clearance −0.74 m), none of them into its three-dimensional volume (smallest three-dimensional clearance 4 mm); with perception, the machine had stopped at other poses, and no path entered its horizontal collision geometry (smallest clearance 3 mm, 33 mm with the blackout). In the 100 trials with perception, the tracker revoked the worker status of no person track (Section 5). Without the monitor, the machine-induced clearance fell to −0.53 m with the barrier rows alone (19 encounters below the standoff) and to −0.77 m without safety layers (31 encounters below the standoff, 12 with machine-induced contact),

Table 8
Worker study (P3 with truck loading, 50 prescribed encounters, A–D inside and E outside the envelope): trials, smallest horizontal clearance between the worker disk and the machine's collision circles (all samples and machine-induced samples), trials with a machine-induced clearance below the standoff, trials with three-dimensional contact (machine-induced / all), the largest machine share of the closing while the worker was inside the standoff, protective stops and delivered volume per trial.

			Smallest clearance [m]		Machine-induced	3-D contact	Approach		Deliv.
Condition	Paths	n	all	machine-ind.	below standoff	mach. / all	[mm]	Stops	[m^3]
monitor, exact states	A–D	40	−0.74	4.39	0	0 / 0	6.5	109	0.00
	E	10	0.88	3.24	0	0 / 0	0.1	26	0.00
monitor, perception	A–D	40	0.00	4.16	0	0 / 0	1.0	281	0.00
	E	10	0.89	3.59	0	0 / 0	0.1	85	0.00
monitor, perception, 1 s blackout	A–D	40	0.03	4.16	0	0 / 0	5.5	298	0.00
	E	10	0.89	3.59	0	0 / 0	14.3	87	0.00
barrier rows only, exact states	A–D	40	−0.60	−0.53	19	0 / 2	2209.6	0	0.49
	E	10	1.05	1.34	2	0 / 0	121.2	0	0.72
no safety layers	A–D	40	−0.78	−0.77	31	12 / 13	3425.9	0	1.00
	E	10	−0.59	−0.59	6	0 / 0	2290.8	0	1.00

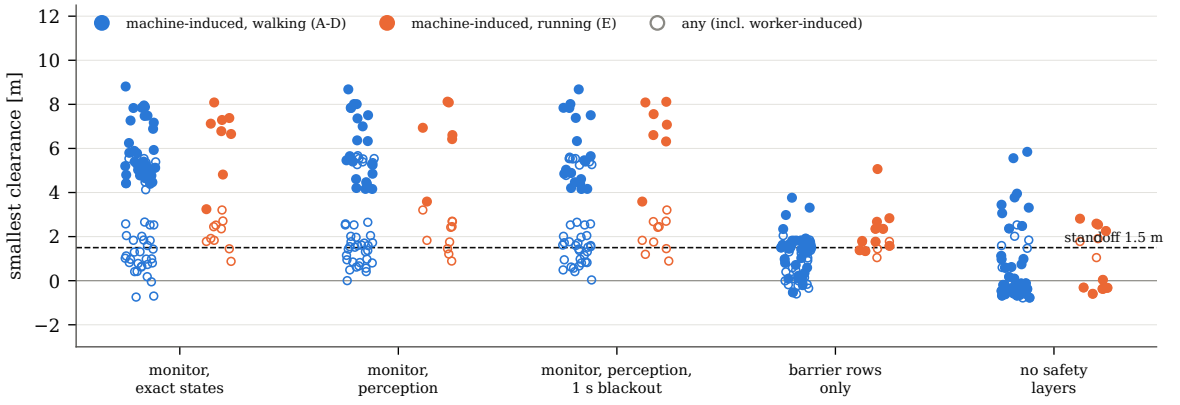

Figure 8: Smallest horizontal clearance between the worker and the machine's collision circles in every encounter under the five conditions: over the machine-induced samples (filled; walking workers of families A–D and running workers of family E) and over all samples, including the closing by the worker's own motion (open), with the standoff (dashed).

so the encounters do expose the worker. Running workers (family E, outside the envelope) came no closer than 3.24 m by the machine's motion with exact states and 3.59 m with perception, although the monitor's bounds assume walking speed; ten encounters per condition do not show that this holds for running persons in general.

10.7. Learned residuals

Each residual was trained from five seeds (3101–3105) with the gate of Section 8.2, and every deployed policy (the last accepted checkpoint) was evaluated once on held-out cohorts of 64 trials in the nominal and stress domains of the reduced-order twin (Table 9). The intervals describe the variation over the five training seeds on these fixed cohorts, not that of the cohorts themselves, and the two cohorts differ both in their plant ranges and in their scenario draws, so that rates are compared within a cohort, between a policy and its prior.

Navigation: the gate verified 12 checkpoints per seed, which arrived at the stance in 72–94 % of the verification trials, and accepted 0–1; four seeds accepted no update and deployed the initial policy, the prior with the residual fitted to zero (which reproduces the prior only approximately), so that the navigation comparison is mostly one between the prior and this copy of it. The prior completed 90.6 % (nominal) and 84.4 % (stress) of the held-out trials and the learned policies 89.1 % and 84.7 % on average (differences −1.6

Table 9
Learned residuals versus their priors on held-out cohorts of the reduced-order twin (64 trials each; five training seeds; mean and 95 % Student-t interval over the seeds; the prior is deterministic on the fixed cohort).

Skill	Cohort	Metric	Prior	Learned [95 % CI]	Learned − prior [95 % CI]
Navigation	nominal	Completion (%)	90.6	89.1 [84.7, 93.4]	−1.6 [−5.9, 2.8]
		Final position error (m)	0.171	0.170 [0.161, 0.180]	−0.001 [−0.010, 0.009]
Navigation	stress	Completion (%)	84.4	84.7 [83.8, 85.6]	0.3 [−0.6, 1.2]
		Final position error (m)	0.180	0.183 [0.175, 0.191]	0.003 [−0.004, 0.011]
Excavation	nominal	Completed cycles (%)	67.2	65.6 [62.9, 68.4]	−1.6 [−4.3, 1.2]
		Delivered (m^3/trial)	0.495	0.485 [0.477, 0.493]	−0.010 [−0.018, −0.001]
		Spilled in carry (m^3/trial)	0.0035	0.0043 [0.0020, 0.0066]	0.0009 [−0.0014, 0.0032]
Excavation	stress	Completed cycles (%)	79.7	81.2 [77.4, 85.1]	1.6 [−2.3, 5.4]
		Delivered (m^3/trial)	0.580	0.610 [0.579, 0.642]	0.031 [−0.0003, 0.062]
		Spilled in carry (m^3/trial)	0.0252	0.0126 [0.0007, 0.0245]	−0.0126 [−0.0245, −0.0007]

[−5.9, +2.8] and +0.3 [−0.6, +1.2] percentage points). Excavation (coverage cycle, inverse valve model for actuation): the gate accepted 0–2 of 20 checkpoints per seed (one seed accepted none); the learned cycles were completed in 65.6 % (nominal) and 81.2 % (stress) of the trials against 67.2 % and 79.7 % for the prior (−1.6 [−4.3, +1.2] and +1.6 [−2.3, +5.4] points); the delivered volume per trial differed by −0.010 [−0.018, −0.001] and +0.031 [−0.0003, +0.062] m^3 (learned minus prior), and the spill in the carry by +0.0009 [−0.0014, +0.0032] and −0.0126 [−0.0245, −0.0007] m^3. No learned policy breached a safety-critical row (self, obstacle, worker, footprint) on the held-out cohorts. Under stress, however, the learned excavation residuals brought the front inside the standoff of a critical row in more trials than the prior (15.4 against 11 of 64; difference [+2.5, +6.3]) and crossed a ground row (front below the admissible floor) in more trials (14.2 against 9; [+3.4, +7.0]); in the nominal domain they did so in 17.0 against 20 trials ([−4.8, −1.2]).

The residuals therefore did not improve on their structured priors in the task: no completion difference is resolved, and the delivered volume is lower in the nominal domain. The resolved gains, a lower carry spill under stress and fewer ground-row crossings in the nominal domain, come with more standoff entries and more ground-row crossings under stress. The gate did what it was designed to do—it kept updates from lowering the task performance on its verification cohort—but the improvements that it admitted on that cohort did not carry over to the held-out cohorts. We did not establish the reason. The priors already complete most trials, and both safety layers reshape every action of a residual, so that an update changes the executed motion less than the policy output; whether a larger budget or a gate on a larger verification cohort would change the outcome was not tested. The evaluated system uses the priors.

11. Discussion

What completes a pit. With exact knowledge of the soil, none of the ablations of Section 10.3 kept a pit from the grade; they changed how long it took and how close its boundary came to the tolerance. The pass margins and the far-wall attitude mattered most for P3: without them, the pit took 76 % and 50 % longer, and with the bucket curled at the far wall it ended at the edge of the tolerance (7.9 cm against 5.1 cm with the design). Without the stable-wall layering, both pits took longer (23 % and 32 %), and the coarse swing band and the fixed dump point changed little in a single pit. In the missions, with the complete design and perception in the loop, the excavation of a pit took 533 s on average, and in the median pit at grade the largest deviation from the design grade was 4.1 cm (Section 10.5). Whether the missions would complete their pits without each element was not tested, and the ablations cover one scenario seed of two pits that were used in the development of the design. The corresponding elements of the task layer resemble familiar earthwork measures (benched walls, an overcut toe, clean-up passes), but in the twin they were derived from the mechanisms that the diagnostics exposed: the failure step of the soil model, the volume that a pass collects relative to the bucket, the delay with which the teeth reach the floor after penetrating, the heel of the curled bucket at the far wall, the spill of the closing bucket and of a bucket opened over rows that the LiDAR does not see, the rate band of the swing's inverse valve map, the stiffness of a buried

bucket against a lateral correction and the standoff that the safety layers keep from objects near the pit, including soil heaped beside it, which the obstacle map takes for an object unless the pit's soil zone covers it. Because these mechanisms belong to the twin, the specific values (0.2 m layers, margins of 0.2 m and 0.5 m, the 0.07 m allowance, the far-wall attitude, the 0.1 m threshold of the thin last pass and the 1.3 m soil zone) are parameters to be identified on a real machine and soil, not general constants; Section 10.4 shows how the completion of a pit with the fixed design depends on the soil parameters.

What completion costs. The missions removed 3.1 m^3 of bank soil per pit on average, 2.0 times the design volume of a pit: the stepped walls of the stable-wall layering and the notches that the pass margins cut beyond the near and far walls are excavated as well. A pass carried 0.44 m^3 of loose soil on average, less than half of the content at which the cycle closes the bucket, because the layers are thin and the finishing passes collect little, and a ten-pit mission took 1.54 h with the slower swing and travel of the twin. The task layer was designed for completion to grade, not for output, and its efficiency was not optimized; thicker layers where the wall admits them, fewer finishing passes and a planned order of the strips would be the first candidates.

What worker protection costs. The identified stopping time of the twin (up to 2.25 s until the upper structure is at rest) makes the walking distance of a person during the stop a large term of the protective separation distance of a moving machine; ahead of a circle that moves at 1 m/s, it is the largest (4.6 m of 8.0 m). A person within 8.0 m of a collision circle that moves toward the person at 1 m/s therefore restricts the machine, and productive work requires a site rule that keeps people farther away (12.5 m in the missions, more than the 12.2 m needed ahead of the fastest measured circle speed, although less than the 13.7 m at the largest effective speed of the measured stops). In the missions, protection cost 4.6 % of the time in protective stops and 7.9 % at limited speed (Section 10.5). It could be reduced by a faster stop (a stiffer swing brake, a shorter anti-reaction window), by bounds that depend on the payload and on the direction of the motion, and by better tracking of people around the machine; none of these is a matter of the monitor's logic. The worker study separates what the monitor provides from what it cannot provide: it limits the clearance that the machine's own motion takes from a person, but a person who walks into a machine at rest or under a raised front reduces the clearance by the person's own motion, and only site rules and exclusion zones prevent this.

Learning. Under the budgets studied, the learned residuals did not improve on the structured priors, and the evaluated system uses the priors. This does not show that learning cannot help an excavator. For skills with strong priors and with safety layers that shape every action, however, the gains that the gate admitted on its verification cohort did not carry over to held-out cohorts, so an improvement has to be confirmed on independent cohorts and on the task-level criteria of the site.

Limitations. All results are obtained in simulation; Table 10 lists what each part of the twin was checked against and what remains unvalidated. The front and its hydraulics are calibrated against field records, but the swing and travel drives are sized from the specification sheet and are slower than the machine (Section 10.1), so swing and travel times are longer than on the machine and the stopping bounds belong to the twin; on a machine they must be identified again with the procedure of Section 6.1. The soil is a heightfield with literature parameters; it does not represent rupture surfaces behind the wall, water, layered or stony ground, and no digging forces were measured on the machine. Its displacement rule moves soil that the bucket shell overlaps to the first cells outside the bucket's footprint along the shell's motion, which can carry a few litres more than a metre, across the footprint; in the sixth earlier run, this left a pit with one cell 1.3 cm outside the tolerance (Section 9), and in the final missions P8 of seed 424 with one cell 0.6 cm outside it (Section 10.5). The bucket model holds more than the machine's bucket (Section 10.1), so the loads of a dig cycle are larger than on the machine. The sensitivity study changes one parameter at a time on one pit and one scenario seed; it shows how sensitive that pit is, not how often pits would fail across soils.

The design parameters were developed on some of the evaluated pits and seeds, whose results are in-sample; of the missions, only those of seeds 421–425 were held out (Section 9), and every earlier held-out set of the missions exposed a gap that the design was then changed to close, so rarer failure modes may remain. The sensors are simulated with range noise and dropouts but without dust, rain or reflections, the camera segmentation is trained on synthetic images, and the georeferenced design grade assumes machine-control positioning without error. The site workers follow scripted routes and site rules, and the workers of the protection study follow prescribed paths; neither represents groups, children, distraction or erratic motion.

Table 10
Basis and checks of the parts of the twin that drive the results. The checks of the soil, the stopping bounds and the integrated workflow are made within the twin and are not validations against the machine.

Part	Basis in the twin	Checked against	Not validated
Front and its hydraulics	closed-chain front, electro-hydraulic circuit	two field sessions in free air (replay errors 0.16° to 0.44° in windows, up to 1.03° over a session; prediction of the other session up to 2.49°; motion onset within 1.9 duty points)	digging loads on the machine
Working ranges, digging forces	calibrated geometry, cylinder areas, relief pressures	specification sheet (within 0.4 % and 2.4 % to 3.9 %)	–
Swing and travel drives	sized from the specification, assumed valve curves	specification speeds (22 % and 13 % slower)	field dynamics and stops
Bucket	convex hull of the shell	rated heaped capacity (46 % larger)	cavity shape, retention
Soil	heightfield, fundamental earthmoving equation, literature parameters	none; sensitivity of one pit to one parameter at a time (Section 10.4)	cutting forces, wall failure, spill
Perception	simulated LiDARs, scanners and cameras; network trained on rendered images	intersection over union on rendered images; detection errors in the missions	dust, rain, reflections, real images
Stopping bounds	controlled stops of the twin	two further sets of stops of the twin; worker study	stops of the machine, travel stops
Integrated workflow	full twin with perception	ten ten-pit missions (five on held-out seeds), single pits, ablations	field trials

The monitor's envelope excludes travel between the pits, for which no stops were identified, and a person who walks back toward the machine while hidden from every sensor (Section 6.3). The tracker's revocation of the worker status (Section 5) takes a person-sized segment that has lain on long-observed static surface for 10 s for structure; a worker hidden behind a structure whose track passes to a person-sized part of that structure, instead of being lost, would lose the status after 10 s, without the reachable set that a lost track receives. The pits are small footings in level prepared ground; trenches, basements, slopes and multiple machines were not studied. The ten ten-pit missions are independent scenario seeds of one site and one work order; the pits of a mission are not independent of each other. The worker-protection rule is a design rule of this simulation study and not a certified safety function; ISO 21815 and ISO/TS 15066 were used as references for the concepts, not as a basis for compliance.

12. Conclusion

SYNEX connects site perception to the valve commands of a 30 t hydraulic excavator and is designed to complete multi-pit work orders to an every-point grade criterion; it was evaluated on a digital twin whose front and its hydraulics are calibrated against field records. Three elements are central. First, a construction-task layer that derives its passes from the mechanisms that decide the boundary of a pit: stable-wall layers below the failure step of the soil model, pass margins for the cutting process, a bucket attitude that keeps the heel clear of the far wall, a thin last pass on every grade strip, a carry to the dump for every grade-layer load, a swing actuation fine enough for the strip lines, dump points within the reach of the curled bucket and full trucks exchanged. With perception in the loop, 99 of the 100 pits of ten ten-pit missions reached the design grade at every point, 49 of the 50 pits of the five held-out missions among them, and nine of the ten missions completed the work order. The pit that missed the criterion was left 0.6 cm outside the tolerance at one cell by the displacement rule of the soil model, which moved soil across the bucket's footprint onto floor that an earlier pass had cut to the grade. The design ablations, on two pits of one scenario seed, indicate what the individual elements contribute to the time and the margin with which a pit is completed. Second,

a control stack whose two control-barrier-function layers, one on the motion reference and one on the valve duties, reduce violations of the modelled boundaries to a small residual despite the dead time, deadzone and shared pump flow of the hydraulics, with a planner that stays synchronized with the measured motion and five liveness rules that release the machine from states in which the safety layers, the tracker and the task layer would otherwise hold it indefinitely. Third, worker protection by a separation monitor on the swept stopping region, whose stopping bounds are evaluated at an effective speed that accounts for the motion reference leading the hydraulics; identified and checked on separate sets of controlled stops, the bounds held for every stop of a third set measured after they were fixed, and in 40 prescribed encounters inside the stated envelope no clearance below the standoff was observed during machine approach faster than $0.05\,\mathrm{m/s}$, and the revocation of the worker status that released the missions from person tracks on static structure did not act in these encounters.

The results are those of a simulation. The front and its hydraulics are calibrated against field records, while the swing and travel drives, the soil and the sensors are modelled from specifications and literature; the margins of the task layer and the stopping bounds are therefore parameters to be identified again on a real machine and soil, with the procedures described here, and the worker-protection rule is a design rule, not a certified safety function. Each earlier set of held-out missions exposed a gap of the design that the development missions had not shown, and five held-out missions cannot exclude rarer failure modes. Learned residuals did not improve on the structured priors under the budgets studied, and the evaluated system uses the priors. Field trials of the task layer and of the identification of the stopping bounds, trenches and slopes, and several machines on one site are the next steps.

CRediT authorship contribution statement

Mehdi Heydari Shahna: Writing – original draft, Validation, Methodology, Investigation, Formal analysis, Software, Data curation, Conceptualization.
Seihun Kim: Writing – review, Data curation.
Soyi Jung: Writing – review & editing, Supervision, Funding acquisition.
Soohyun Park: Writing – review & editing, Supervision, Funding acquisition.
Jouni Mattila: Writing – review & editing, Resources, Supervision.
Joongheon Kim: Writing – review & editing, Supervision, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This work was supported by the Korean Ministry of Trade, Industry and Energy (MOTIE) under Grant RS-2024-00442168, through the Mechanical Equipment Industry Technology Development Program administered by KEIT.

Data availability

All code, trained policies, configuration files, result files and mission logs used in this study, together with a supplementary video, are provided as supplementary material with the manuscript submission. After acceptance, all of it will be made publicly available in an open-source GitHub repository, whose link will be given in the published article. Raw field-calibration records are not included; the supplied parameter and validation summaries do not permit independent recalibration.

References

[1] O. M. U. Eraliev, K.-H. Lee, D.-Y. Shin, C.-H. Lee, Sensing, perception, decision, planning and action of autonomous excavators, Automation in Construction 141 (2022) 104428. doi:10.1016/j.autcon.2022.104428.

[2] J. Feng, W. Chen, C. Liu, X. Liu, B. Cao, F. Xu, Advancements and challenges in earthmoving equipment automation using core technologies, Automation in Construction 179 (2025) 106472. doi:10.1016/j.autcon.2025.106472.

[3] J. Mattila, J. Koivumäki, D. G. Caldwell, C. Semini, A survey on control of hydraulic robotic manipulators with projection to future trends, IEEE/ASME Transactions on Mechatronics 22 (2) (2017) 669–680. doi:10.1109/TMECH.2017.2668604.

[4] D. Jud, S. Kerscher, M. Wermelinger, E. Jelavic, P. Egli, P. Leemann, G. Hottiger, M. Hutter, HEAP - the autonomous walking excavator, Automation in Construction 129, art. no. 103783 (Sep. 2021). doi:10.1016/j.autcon.2021.103783.

[5] L. Zhang, J. Zhao, P. Long, L. Wang, L. Qian, F. Lu, X. Song, D. Manocha, An autonomous excavator system for material loading tasks, Science Robotics 6 (55), art. no. eabc3164 (Jun. 2021). doi:10.1126/scirobotics.abc3164.

[6] J. Seo, S. Lee, J. Kim, S. K. Kim, Task planner design for an automated excavation system, Automation in Construction 20 (7) (2011) 954–966. doi:10.1016/j.autcon.2011.03.013.

[7] D. Jud, G. Hottiger, P. Leemann, M. Hutter, Planning and control for autonomous excavation, IEEE Robotics and Automation Letters 2 (4) (2017) 2151–2158. doi:10.1109/LRA.2017.2721551.

[8] P. Egli, M. Hutter, A general approach for the automation of hydraulic excavator arms using reinforcement learning, IEEE Robotics and Automation Letters 7 (2) (2022) 5679–5686. doi:10.1109/LRA.2022.3152865.

[9] P. Egli, D. Gaschen, S. Kerscher, D. Jud, M. Hutter, Soil-adaptive excavation using reinforcement learning, IEEE Robotics and Automation Letters 7 (4) (2022) 9778–9785. doi:10.1109/LRA.2022.3189834.

[10] J. Cho, M. Shin, J. Kim, S. Jung, Collaborative learning architecture for autonomous excavator planning and execution, Automation in Construction 182 (2026) 106742. doi:10.1016/j.autcon.2025.106742.

[11] A. Molaei, M. Heravi, R. Ghabcheloo, Data-driven rock capture with a standard excavator via model-free reinforcement learning, Automation in Construction 185 (2026) 106869. doi:10.1016/j.autcon.2026.106869.

[12] A. D. Ames, X. Xu, J. W. Grizzle, P. Tabuada, Control barrier function based quadratic programs for safety critical systems, IEEE Transactions on Automatic Control 62 (8) (2017) 3861–3876. doi:10.1109/TAC.2016.2638961.

[13] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, P. Tabuada, Control barrier functions: Theory and applications, in: 2019 18th European Control Conference (ECC), IEEE, Naples, Italy, 2019, pp. 3420–3431. doi:10.23919/ECC.2019.8796030.

[14] Q. Nguyen, K. Sreenath, Exponential control barrier functions for enforcing high relative-degree safety-critical constraints, in: 2016 American Control Conference (ACC), IEEE, Boston, MA, USA, 2016, pp. 322–328. doi:10.1109/ACC.2016.7524935.

[15] W. Xiao, C. Belta, Control barrier functions for systems with high relative degree, in: 2019 IEEE 58th Conference on Decision and Control (CDC), IEEE, Nice, France, 2019, pp. 474–479. doi:10.1109/CDC40024.2019.9029455.

[16] T. G. Molnar, A. K. Kiss, A. D. Ames, G. Orosz, Safety-critical control with input delay in dynamic environment, IEEE Transactions on Control Systems Technology 31 (4) (2023) 1507–1520. doi:10.1109/TCST.2022.3227451.

[17] M. H. Shahna, J. Mattila, Tracking limitations and rate-consistent safety filtering for constrained mechanical systems, Research Square preprint (Oct. 2026). doi:10.21203/rs.3.rs-11381118/v1.
URL https://www.researchsquare.com/article/rs-11381118/v1

[18] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, U. Topcu, Safe reinforcement learning via shielding, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 32, 2018. doi:10.1609/aaai.v32i1.11797.

[19] R. Cheng, G. Orosz, R. M. Murray, J. W. Burdick, End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33, 2019, pp. 3387–3395. doi:10.1609/aaai.v33i01.33013387.

[20] International Organization for Standardization, ISO/TS 15066:2016 robots and robotic devices – collaborative robots, Technical specification, ISO, Geneva (2016).

[21] International Organization for Standardization, ISO 13855:2024 safety of machinery – positioning of safeguards with respect to the approach of the human body, International standard, ISO, Geneva (2024).

[22] International Organization for Standardization, ISO 21815-1:2022 earth-moving machinery – collision warning and avoidance – part 1: General requirements, International standard, ISO, Geneva (2022).

[23] J. Teizer, B. S. Allread, C. E. Fullerton, J. Hinze, Autonomous pro-active real-time construction worker and equipment operator proximity safety alert system, Automation in Construction 19 (5) (2010) 630–640. doi:10.1016/j.autcon.2010.02.009.

[24] B.-W. Jo, Y.-S. Lee, J.-H. Kim, D.-K. Kim, P.-H. Choi, Proximity warning and excavator control system for prevention of collision accidents, Sustainability 9 (8) (2017) 1488. doi:10.3390/su9081488.

[25] X. Lin, Z. Guo, X. Jin, H. Guo, Digital twin-enabled safety monitoring system for seamless worker-robot collaboration in construction, Automation in Construction 174 (2025) 106147. doi:10.1016/j.autcon.2025.106147.

[26] T. Fu, T. Zhang, Y. Lv, X. Song, G. Li, H. Yue, Digital twin-based excavation trajectory generation of uncrewed excavators for autonomous mining, Automation in Construction 151 (2023) 104855. doi:10.1016/j.autcon.2023.104855.

[27] M. H. Shahna, J. Kim, J. Mattila, RoboCompiler: Graph-native compilation of closed-chain robots for consistent modeling, control, and simulation, arXiv preprint arXiv:2609.35717 (Sep. 2026). arXiv:2609.35717, doi:10.48550/arXiv.2609.35717.
URL https://arxiv.org/abs/2609.35717

[28] M. H. Shahna, S. Kim, S. Jung, S. Park, J. Mattila, J. Kim, Graph-based modeling and simulation of a hydraulic excavator using RoboCompiler, Preprints.org preprint (Oct. 2026). doi:10.20944/preprints202610.1019.v1.
URL https://www.preprints.org/manuscript/202610.1019/v1

[29] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, G. State, Isaac Gym: High performance GPU based physics simulation for robot learning, in: Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks, Vol. 1, 2021.

[30] N. Rudin, D. Hoeller, P. Reist, M. Hutter, Learning to walk in minutes using massively parallel deep reinforcement learning, in: Proceedings of the 5th Conference on Robot Learning, Vol. 164 of Proceedings of Machine Learning Research, PMLR, 2022, pp. 91–100.

[31] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, arXiv preprint arXiv:1707.06347 (Jul. 2017). arXiv:1707.06347, doi:10.48550/arXiv.1707.06347.

[32] J. Schulman, P. Moritz, S. Levine, M. Jordan, P. Abbeel, High-dimensional continuous control using generalized advantage estimation, in: International Conference on Learning Representations (ICLR), 2016. arXiv:1506.02438.

[33] D. Chen, B. Zhou, V. Koltun, P. Krähenbühl, Learning by cheating, in: Proceedings of the Conference on Robot Learning, Vol. 100 of Proceedings of Machine Learning Research, PMLR, 2020, pp. 66–75.

[34] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, M. Hutter, Learning quadrupedal locomotion over challenging terrain, Science Robotics 5 (47), art. no. eabc5986 (Oct. 2020). doi:10.1126/scirobotics.abc5986.

[35] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, M. Hutter, Learning robust perceptive locomotion for quadrupedal robots in the wild, Science Robotics 7 (62), art. no. eabk2822 (Jan. 2022). doi:10.1126/scirobotics.abk2822.

[36] T. Johannink, S. Bahl, A. Nair, J. Luo, A. Kumar, M. Loskyll, J. A. Ojea, E. Solowjow, S. Levine, Residual reinforcement learning for robot control, in: 2019 International Conference on Robotics and Automation (ICRA), IEEE, Montreal, QC, Canada, 2019, pp. 6023–6029. doi:10.1109/ICRA.2019.8794127.

[37] T. Silver, K. Allen, J. Tenenbaum, L. Kaelbling, Residual policy learning, arXiv preprint arXiv:1812.06298 (Dec. 2018). arXiv:1812.06298, doi:10.48550/arXiv.1812.06298.

[38] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, J. Nieto, Voxblox: Incremental 3D Euclidean signed distance fields for on-board MAV planning, in: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, Vancouver, BC, Canada, 2017, pp. 1366–1373. doi:10.1109/IROS.2017.8202315.

[39] A. Millane, H. Oleynikova, E. Wirbel, R. Steiner, V. Ramasamy, D. Tingdahl, R. Siegwart, nvblox: GPU-accelerated incremental signed distance field mapping, in: 2024 IEEE International Conference on Robotics and Automation (ICRA), IEEE, Yokohama, Japan, 2024, pp. 2698–2705. doi:10.1109/ICRA57147.2024.10611532.

[40] I. Bogoslavskyi, C. Stachniss, Fast range image-based segmentation of sparse 3D laser scans for online operation, in: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, Daejeon, South Korea, 2016, pp. 163–169. doi:10.1109/IROS.2016.7759050.

[41] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Vol. 9351 of Lecture Notes in Computer Science, Springer International Publishing, 2015, pp. 234–241. doi:10.1007/978-3-319-24574-4_28.

[42] A. R. Reece, The fundamental equation of earth-moving mechanics, Proceedings of the Institution of Mechanical Engineers, Conference Proceedings 179 (6) (1964) 16–22. doi:10.1243/PIME_CONF_1964_179_134_02.

[43] E. McKyes, Soil Cutting and Tillage, Vol. 7 of Developments in Agricultural Engineering, Elsevier, Amsterdam, The Netherlands, 1985. doi:10.1016/C2009-0-01148-9.

[44] M. Servin, T. Berglund, S. Nystedt, A multiscale model of terrain dynamics for real-time earthmoving simulation, Advanced Modeling and Simulation in Engineering Sciences 8 (1), art. no. 11 (2021). doi:10.1186/s40323-021-00196-3.

[45] R. W. Sumner, J. F. O'Brien, J. K. Hodgins, Animating sand, mud, and snow, Computer Graphics Forum 18 (1) (1999) 17–26. doi:10.1111/1467-8659.00299.

[46] E. Todorov, T. Erez, Y. Tassa, MuJoCo: A physics engine for model-based control, in: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, Vilamoura-Algarve, Portugal, 2012, pp. 5026–5033. doi:10.1109/IROS.2012.6386109.

[47] T. F. Chan, R. V. Dubey, A weighted least-norm solution based scheme for avoiding joint limits for redundant joint manipulators, IEEE Transactions on Robotics and Automation 11 (2) (1995) 286–292. doi:10.1109/70.370511.

[48] D. Goldfarb, A. Idnani, A numerically stable dual method for solving strictly convex quadratic programs, Mathematical Programming 27 (1) (1983) 1–33. doi:10.1007/BF02591962.

[49] International Electrotechnical Commission, IEC 61496-3:2018 safety of machinery – electro-sensitive protective equipment – part 3: Particular requirements for active opto-electronic protective devices responsive to diffuse reflection (AOPDDRs), International standard, IEC, Geneva (2018).

[50] Doosan Infracore, DX300LC-5 specification sheet, rev. 02 (2014-11-18), https://eu.develon-ce.com/media/download/6429 (2014).

[51] International Organization for Standardization, ISO 6015:2006 earth-moving machinery – hydraulic excavators and backhoe loaders – methods of determining tool forces, International standard, ISO, Geneva (2006).

[52] International Organization for Standardization, ISO 7451:2007 earth-moving machinery – volumetric ratings for hoe-type and grab-type buckets of hydraulic excavators and backhoe loaders, International standard, ISO, Geneva (2007).



Mehdi Heydari Shahna received the D.Sc. (Tech.) degree with distinction in Automation Science and Engineering from Tampere University, Finland, in 2025. He is currently a postdoctoral researcher at Tampere University. His research interests include physical AI, robot learning, safe control, and robust control.



Seihun Kim is currently pursuing the M.S. degree in defense convergence technology at Korea University, Seoul, Republic of Korea. He received the B.S. degree in electronic engineering from Soongsil University, Seoul, Republic of Korea, in 2022. His research interests include deep learning algorithms and their applications to mobility and networking.

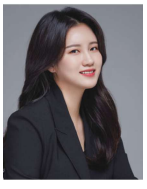

Soyi Jung (Senior Member, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical and computer engineering from Ajou University, Suwon, South Korea, in 2013, 2015, and 2021, respectively. She has been an Associate Professor with the Department of Electrical and Computer Engineering, Ajou University, since September 2025.



Soohyun Park (Member, IEEE) received the B.S. degree in computer science and engineering from Chung-Ang University, Seoul, South Korea, in February 2019, and the Ph.D. degree in electrical and computer engineering from Korea University, Seoul, in August 2023. She has been an Assistant Professor with the Division of Computer Science, Sookmyung Women's University, Seoul, since March 2025.



Jouni Mattila received a Ph.D. in automation engineering from Tampere University of Technology, Tampere, Finland, in 2000. He is currently a professor of machine automation with the Unit of Automation Technology and Mechanical Engineering at Tampere University. His research interests include machine automation, nonlinear-model-based control of robotic manipulators, and energy-efficient control of heavy-duty mobile manipulators.



Joongheon Kim (Senior Member, IEEE) received the B.S. and M.S. degrees in computer science and engineering from Korea University, Seoul, South Korea, in 2004 and 2006, respectively, and the Ph.D. degree in computer science from the University of Southern California, Los Angeles, CA, USA, in 2014. He has been a Professor with the Department of Electrical and Computer Engineering, Korea University, since September 2019.