*Human Factors in Design, Engineering, and Computing*, Vol. 238, 2026, 1306–1315
https://doi.org/10.54941/ahfe1008359

# Autonomous Hospitality Yachts: Designing Trustworthy Human-AI Experiences in Future Maritime Environments

**Laura Pagani, Paolo Gemelli, Alessandro Bertirotti, and Mario Ivan Zignego**

University of Genoa, Genoa, Italy

## ABSTRACT

The development of Maritime Autonomous Surface Ships (MASS), artificial intelligence and cyber-physical systems is reshaping future marine environments. While maritime autonomy has mainly addressed cargo transport, surveillance and industrial operations, the leisure marine sector remains underexplored. This paper proposes a human-centred conceptual framework for Autonomous Hospitality Yachts (AHY), understood as supervised intelligent vessels integrating marine autonomy, AI-driven hospitality services, remote oversight and adaptive human-AI interaction. The framework combines a layered system architecture with four human-factors requirements - perceived safety, trust in AI behaviour, emotional comfort and perceived operational reliability - grounded in neuroscientific and cognitive models of guest experience. Rather than advocating fully unsupervised yachts, the paper argues for progressive, domain-bounded and explainable autonomy aligned with MASS principles and Operational Design Domains. The framework is intended as a research and design model requiring future empirical validation, not as a validated deployment architecture.

**Keywords:** Autonomous hospitality yachts, Human-centred AI, Maritime autonomy, Trustworthy AI, Human factors, Future of work, Perceived safety, Supervised autonomy

## INTRODUCTION

Artificial intelligence, sensor fusion, digital twins and cyber-physical systems are progressively transforming marine operations. In the context of Maritime Autonomous Surface Ships (MASS), current research and regulatory discussions primarily address cargo transport, naval applications, port operations and industrial logistics. By contrast, the impact of autonomy on the leisure marine sector remains underexplored, although yachts are complex inhabited environments in which navigation, hospitality, safety, comfort and emotional experience are tightly interwoven.

This paper introduces the concept of the Autonomous Hospitality Yacht (AHY) as an emerging human-centred marine system. An AHY is not conceived as a fully unmanned luxury vessel, but as a supervised intelligent hospitality environment in which autonomous navigation functions, cyber-physical vessel monitoring, AI-supported services, remote oversight and

Received June 27, 2026; Revised September 22, 2026; Accepted October 3, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

adaptive human-AI interaction are integrated within a coherent design framework. Autonomy in leisure environments should not be evaluated solely through technical performance metrics: guests, crew and remote operators should be able to understand, trust and meaningfully supervise the system.

The paper therefore follows a guiding principle: in autonomous hospitality, autonomy need not be constantly visible, but care must be perceptible through comfort, safety, attentiveness and meaningful human control.

The contribution is threefold. First, the paper defines AHYs as supervised cyber-physical-service systems. Second, it proposes a layered architecture integrating marine autonomy, vessel management, AI-driven hospitality, human-AI interaction and remote safety governance. Third, it discusses future-of-work implications, where crew roles shift from direct execution toward supervision, exception management, service orchestration and trust calibration.

## BACKGROUND AND SCOPE

Maritime autonomy is generally described as a spectrum of functions with different levels of human involvement. IMO discussions on MASS distinguish ships with automated decision support, remotely controlled ships with seafarers on board, remotely controlled ships without seafarers, and fully autonomous ships (International Maritime Organization, 2021). For yachts carrying guests, autonomy should be linked to specific Operational Design Domains (ODD), environmental conditions, regulatory constraints and human oversight mechanisms, not treated as an absolute property.

Navigation autonomy must remain compatible with collision-avoidance responsibilities under COLREGs (International Maritime Organization, 1972). Safety analyses of unmanned maritime operations further show that autonomy introduces new risk patterns, including remote-control dependencies, cyber-physical vulnerabilities and degraded human situation awareness (Wrobel et al., 2017). These issues are particularly acute in hospitality yachts, where guests are not domain experts and interpret system behaviour through perceived safety rather than technical knowledge.

The framework proposes an extension of the conventional ODD into a Hospitality Operational Design Domain (H-ODD), where weather, traffic and technical limits are considered together with guest comfort tolerance, privacy preferences, acceptable automation visibility, crew availability and remote-supervision readiness.

A critical distinction applies here: the guest profile of a luxury yacht differs fundamentally from that of a passenger-ferry traveller. Ahola et al. (2014) demonstrate that safety perception on passenger ships is shaped by architecture, communication, crew presence and the broader emotional atmosphere on board. On a yacht, these factors are amplified by the intimacy of the environment and by high experiential expectations. Autonomous systems should therefore address a higher bar of contextual, affective and relational communication, not merely technical transparency.

The human-factors literature reinforces this view. Levels of automation should be selected considering not only technical feasibility, but also consequences for human coordination, workload and responsibility

(Parasuraman et al., 2000). Situation awareness is essential for operators supervising dynamic systems (Endsley, 1995), while trust in automation should be calibrated rather than simply maximized (Lee and See, 2004). Appropriate reliance - neither over-trusting nor rejecting automation - is the design target.

## RESEARCH APPROACH

This paper adopts a conceptual framework approach, integrating four knowledge domains: maritime autonomous systems, human-centred AI, yacht hospitality design and human factors in automation. The objective is not to present a final engineering specification, but to organize design dimensions that should guide future prototypes, simulations and empirical studies.
The framework rests on three foundational assumptions:

- autonomy in leisure marine environments should be progressive, supervised and domain-bounded;
- hospitality autonomy must support perceived safety, emotional comfort, privacy-aware personalization and explainable interaction;
- the future of work on board should be reinterpreted rather than erased, because autonomous hospitality redistributes work among onboard personnel, remote supervision centres and AI systems.

## DEFINITION AND LAYERED ARCHITECTURE

An Autonomous Hospitality Yacht can be defined as a supervised intelligent marine platform that integrates autonomous navigation functions, cyber-physical vessel monitoring, AI-driven hospitality services and adaptive human-AI interaction to support safer, more personalized and more trustworthy onboard experiences. Five elements are central: the vessel remains a regulated marine system; it operates as a cyber-physical system; it functions as a hospitality environment; it requires user-specific explanation strategies; and it remains accountable through human oversight.

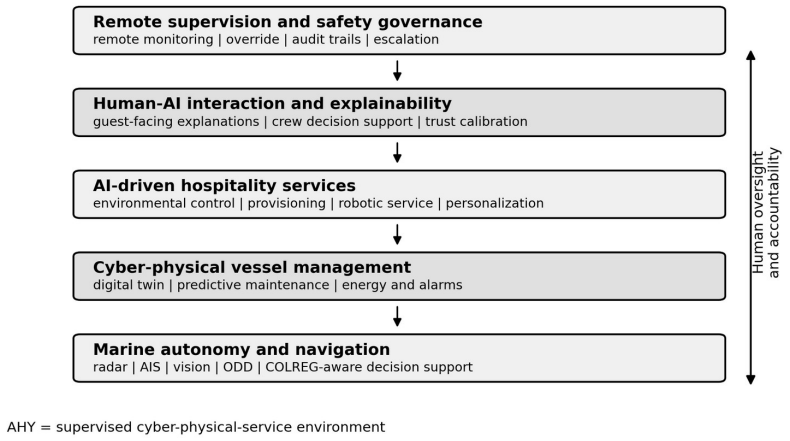



**Figure 1:** Layered AHY architecture. The diagram highlights the dependency between marine autonomy, vessel management, hospitality services, explainable interaction and remote governance. Increasing automation requires stronger human oversight and accountability, not their removal.

**Table 1:** Core layers of the proposed autonomous hospitality yacht framework.

| Layer | Main Functions | Human-Factors Risk | Design Requirement |
|---|---|---|---|
| Marine autonomy and navigation | Radar, AIS, GNSS, vision, route planning, collision-risk support | Opacity of autonomous manoeuvres; overreliance or rejection | ODD visibility, COLREG-aware logic, human override |
| Cyber-physical vessel management | Digital twin, predictive maintenance, energy, alarms, redundancy | Hidden technical degradation affecting comfort and safety | Actionable diagnostics, severity levels, graceful degradation |
| AI-driven hospitality services | Environmental control, provisioning, robotic service, personalization | Intrusive personalization; privacy concerns | Visible and controllable preference learning |
| Human-AI interaction | Guest explanations, crew decision support, remote dashboards | Misunderstanding of system status and limits | Contextual explanations differentiated by user role |
| Remote safety governance | Remote monitoring, escalation, audit trails, cybersecurity | Loss of accountability and delayed intervention | Clear authority, communication channels and procedures |

## HUMAN-FACTORS REQUIREMENTS FOR TRUSTWORTHY AHY

The human-centred quality of an AHY depends on four interdependent requirements: perceived safety, trust in AI behaviour, emotional comfort and perceived operational reliability. These dimensions connect technical autonomy to the lived experience of guests and to the supervisory work of crew and remote operators.

Perceived safety. Perceived safety is not identical to objective safety. In passenger-ship contexts, safety perception is shaped by architecture, crew visibility, communication quality, emotional tone and the presence of other people, not by technical statistics alone (Ahola et al., 2014). Barrett's constructivist model of emotion (Barrett, 2017) and Friston's predictive-processing account (Friston, 2010) offer a theoretical grounding for this claim: safety may be understood as an affective construction generated from bodily signals, prior expectations and prediction-error reduction. AHY design should therefore aim to produce coherent environmental signals - stability, lighting, acoustic quality, tone of communication and visible supervision cues - that support perceived safety and activate positive somatic markers in guests (Damasio, 1994).

Trust in AI behaviour. Trust depends on competence, benevolence and integrity (McKnight et al., 2011). Competence is demonstrated through reliable performance, benevolence through personalization that visibly serves the guest's interests, and integrity through transparent data practices. Critically, trust should be calibrated rather than maximized (Lee and See, 2004). For AHYs, calibration can be supported through explainability: the system should communicate what it is doing, why it is doing it and what will

happen next in language accessible to non-expert guests. Explainability here means contextual, affectively appropriate communication, not algorithm disclosure (Floridi, 2023).

Emotional comfort. Emotional comfort concerns the affective quality of the onboard experience as a whole. It should be pursued through continuous environmental regulation - temperature, humidity, lighting, acoustic quality and vessel stability - before any explicit service interaction. A further dimension concerns anticipation: Lieberman and Long (2018) distinguish dopaminergic wanting, oriented toward possibility, novelty and future improvement, from here-and-now sensory satisfaction. In predictive-processing terms, delight can be interpreted as a positively valenced prediction-error event: the environment confirms safety while exceeding calibrated expectations. This supports a design hypothesis of subtle positive surprise rather than uniform automation.

For example, after dinner an AHY should not silently change lighting and music because it has inferred a preference. It should propose a "quiet evening mode" and allow guests to confirm, modify or reject it, making personalization visible, controllable and emotionally comfortable.

**Perceived operational reliability.** Perceived operational reliability is the experiential assessment guests form of the system's capacity to keep its promises over time. Expectation-confirmation theory shows that satisfaction depends on the gap between expectations and perceived performance (Oliver, 1980; Bhattacherjee, 2001). The design implication is expectation calibration: the system should communicate capabilities realistically, avoid over-promising and then exceed calibrated expectations through personalization, anticipation and graceful degradation when functions are partially unavailable.

**Table 2:** Human-factors requirements, design strategies and evaluation logic.

| Requirement | Meaning in AHY Context | Design Strategy | Possible Evaluation |
|---|---|---|---|
| Perceived safety | Felt sense of protection through environmental and communicative signals | Anticipatory communication, somatic-signal design, visible supervision cues | Perceived-safety interviews, HRV and stress indicators |
| Trust in AI behaviour | System seen as competent, benevolent and consistent; trust calibrated not maximized | Capability/limit communication, contextual explanations, privacy transparency | Trust-in-automation scales, reliance behaviour, qualitative feedback |
| Emotional comfort | Environment feels attentive, anticipatory, non-intrusive and personally responsive | Environmental regulation, designed positive surprise, visible preference control | Experience sampling, comfort ratings, GSR and skin temperature |
| Perceived operational reliability | System maintains its promises; expectations are calibrated and then exceeded | Expectation management, graceful degradation, error recovery communication | Technology acceptance and expectation-confirmation measures |

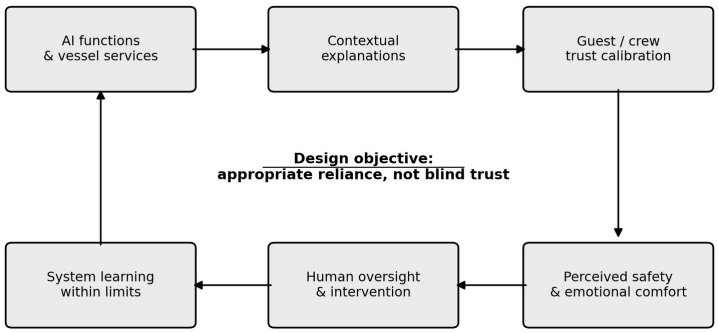



**Figure 2:** Trust-calibration loop for AHYs. AI functions affect guest and crew trust through contextual explanations. Perceived safety and emotional comfort feed human oversight and bounded system learning. The objective is appropriate reliance rather than blind trust.

## FUTURE OF WORK AND PROGRESSIVE ROADMAP

The development of AHYs has direct implications for the future of maritime work. Conventional yacht crews perform navigation support, maintenance, cleaning, provisioning, guest assistance, safety procedures and emergency management. In AHYs, some of these tasks may become partially automated or AI-supported, but this does not imply the disappearance of human work. Work is redistributed across onboard crew, remote operators and intelligent systems.

Crew members may shift from manual execution toward supervisory control, guest-experience orchestration, exception handling, system interpretation and safety communication. Remote operators may become part of the extended crew, supporting navigation, diagnostics, cybersecurity and emergency response from shore-based centres. Future yacht professionals will require competencies in human-AI collaboration, sensor-based diagnostics, explanation of system behaviour, privacy communication and automation-bias management.

The crew of the future thus becomes a trust mediator. When AI recommends delaying an anchorage because of weather evolution, the crew member translates the technical recommendation into a reassuring hospitality decision, preserving both operational safety and the guests' sense of care.

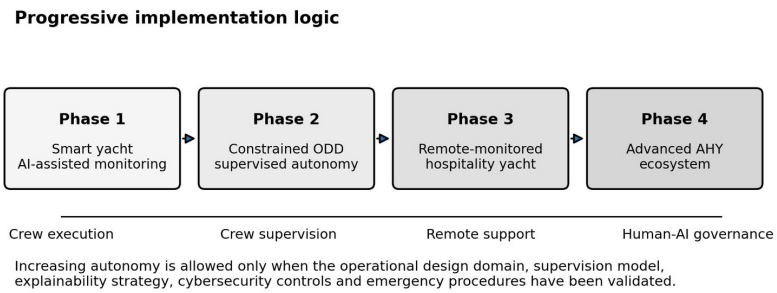



**Figure 3:** Progressive implementation logic. Autonomy should increase only after ODD boundaries, supervision models, explainability strategies, cybersecurity controls and emergency procedures have been validated. The human role evolves from crew execution to human-AI governance.

**Table 3:** Progressive implementation roadmap and validation focus.

| Phase | Technical Scope | Human Role | Validation Priority |
|---|---|---|---|
| 1. Smart yacht | AI-assisted monitoring, predictive maintenance, preference support | Crew operates; AI supports | Usability, acceptance, reliability of alerts |
| 2. Constrained autonomy | Assisted docking, route optimization, collision-risk alerts, adaptive energy | Crew supervises and intervenes | ODD limits, situation awareness, handover procedures |
| 3. Remote-monitored yacht | Remote monitoring centre and diagnostic support | Crew and remote operators coordinate | Latency, authority, emergency escalation protocols |
| 4. AHY ecosystem | High-level supervised autonomy, digital twin, robotic services, adaptive hospitality | Human-AI governance, exception management, service orchestration | Trust calibration against comfort and reliability measures; resilience and accountability audits |

## DESIGN AND EVALUATION IMPLICATIONS

The proposed framework reframes the autonomous yacht as a human-centred intelligent environment rather than a purely technical autonomous vessel. Four design implications follow. First, technical reliability is necessary but insufficient: system behaviour must be understandable and supervisable. Graceful degradation - the capacity to reduce performance in an ordered, communicated way rather than failing abruptly - is as important as uptime. Second, safety design and service design must be integrated. Environmental comfort, emotional acceptance and perceived safety are not separable in a leisure context: the physical design of the vessel, the acoustic environment, lighting and AI communication style are all components of the safety architecture.

In AHYs, hospitality itself becomes the interface: guests interact not only with screens or voice assistants, but with a coordinated environment of light, sound, motion, timing, explanation and human mediation.

Third, following established interaction design principles (Norman, 2013), explanations should be differentiated by user role. Guests require reassurance and control; crew require operational meaning and exception-management support; remote supervisors require diagnostic depth and audit capability. Fourth, privacy-aware personalization must remain visible and controllable. The system should communicate its inferences to guests in ways that invite confirmation or correction, building trust through demonstrated attentiveness rather than silent adaptation.

These evaluation methods should be understood as proposed validation directions, not as results already achieved by the present study. Relevant future measures may include trust-in-automation scales (Jian et al., 2000), technology acceptance instruments (Davis, 1989), expectation-confirmation

measures (Bhattacherjee, 2001), perceived-safety interviews, Heart Rate Variability as a physiological index of stress regulation, galvanic skin response for emotional arousal and scenario interviews covering routine navigation, degraded automation, remote-supervision handover and emergency communication.

## LIMITATIONS AND FUTURE VALIDATION

This paper has limitations that should be explicitly acknowledged. The proposed AHY framework is conceptual and design-oriented. It does not report empirical data from onboard trials, simulator experiments or user studies with yacht guests, crew members or remote supervisors. The framework should therefore be interpreted as a structured research and design proposal rather than as a validated operational model.

The paper also does not provide a complete engineering specification for autonomous yacht deployment. The layered architecture identifies relevant functional domains and human-factors requirements, but it does not define certified control algorithms, detailed cybersecurity architectures, classification procedures or regulatory approval pathways. For this reason, the proposed model should not be read as evidence that highly autonomous hospitality yachts are ready for operational use.

The human-factors claims are based on theoretical integration from maritime autonomy, trust in automation, perceived safety, interaction design and cognitive neuroscience. These literatures provide a coherent basis for hypothesis generation, but their application to luxury yacht environments requires staged empirical validation across guest profiles, cultural expectations, voyage conditions and levels of crew visibility. Future work should therefore move from conceptual modelling to digital-twin scenarios, immersive simulations, onboard prototypes and crew-operator studies before higher levels of AHY autonomy are considered within clearly defined Operational Design Domains.

## CONCLUSION

Autonomous Hospitality Yachts represent a possible future direction for the leisure marine sector, where maritime autonomy, AI-driven services and human-centred hospitality converge. Their value will not depend on autonomous navigation or technical automation alone, but on the capacity to create safe, understandable and emotionally comfortable human-AI experiences.

This paper proposed a conceptual framework for AHYs grounded in supervised autonomy, layered system architecture, explainable interaction, perceived safety, emotional comfort and remote safety governance. The neuroscientific framing - predictive processing, somatic markers and dopaminergic anticipation - provides a theoretically grounded basis for design hypotheses that go beyond usability and address the lived quality of the onboard experience. Autonomous hospitality may also transform

maritime work by shifting crew roles toward supervision, interpretation, service orchestration and exception management.

By bridging naval design, human factors, AI ethics and service engineering, the framework contributes to future research on trustworthy intelligent marine hospitality environments and identifies directions for the design and evaluation of autonomy in leisure maritime contexts. The future of autonomous yachting should not be defined by the absence of humans, but by the quality of the relationships designed among humans, intelligent systems and the sea.

## REFERENCES

Ahola, M., Murto, P., Kujala, P. and Pitkanen, J. (2014) "Perceiving safety in passenger ships - user studies in an authentic environment", Safety Science, 70, pp. 222–232. https://doi.org/10.1016/j.ssci.2014.05.017

Barrett, L.F. (2017) How Emotions Are Made: The Secret Life of the Brain. New York: Houghton Mifflin Harcourt.

Bhattacherjee, A. (2001) "Understanding information systems continuance: An expectation-confirmation model", MIS Quarterly, 25(3), pp. 351–370. https://doi.org/10.2307/3250921

Damasio, A.R. (1994) Descartes' Error: Emotion, Reason, and the Human Brain. New York: Putnam. Davis, F.D. (1989) "Perceived usefulness, perceived ease of use, and user acceptance of information technology", MIS Quarterly, 13(3), pp. 319–340. https://doi.org/10.2307/249008

Endsley, M.R. (1995) "Toward a theory of situation awareness in dynamic systems", Human Factors, 37(1), pp. 32–64. https://doi.org/10.1518/001872095779049543

Floridi, L. (2023) The Ethics of Artificial Intelligence: Principles, Challenges, and Opportunities. Oxford: Oxford University Press.

Friston, K.J. (2010) "The free-energy principle: a unified brain theory?", Nature Reviews Neuroscience, 11(2), pp. 127–138. https://doi.org/10.1038/nrn2787

International Maritime Organization (1972) Convention on the International Regulations for Preventing Collisions at Sea, 1972 (COLREGs). London: IMO.

International Maritime Organization (2021) MSC.1/Circ.1638 - Outcome of the Regulatory Scoping Exercise for the Use of Maritime Autonomous Surface Ships (MASS). London: IMO.

Jian, J.-Y., Bisantz, A.M. and Drury, C.G. (2000) "Foundations for an empirically determined scale of trust in automated systems", International Journal of Cognitive Ergonomics, 4(1), pp. 53–71. https://doi.org/10.1207/S15327566IJCE0401_04

Lee, J.D. and See, K.A. (2004) "Trust in automation: Designing for appropriate reliance", Human Factors, 46(1), pp. 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Lieberman, D.Z. and Long, M.E. (2018) The Molecule of More. Dallas: BenBella Books.

McKnight, D.H., Carter, M., Thatcher, J.B. and Clay, P.F. (2011) "Trust in a specific technology: An investigation of its components and measures", ACM Transactions on Management Information Systems, 2(2), Article 12. https://doi.org/10.1145/1985347.1985353

Norman, D.A. (2013) The Design of Everyday Things, Revised and Expanded Edition. New York: Basic Books.

Oliver, R.L. (1980) "A cognitive model of the antecedents and consequences of satisfaction decisions", Journal of Marketing Research, 17(4), pp. 460–469. https://doi.org/10.1177/002224378001700405

Parasuraman, R., Sheridan, T.B. and Wickens, C.D. (2000) "A model for types and levels of human interaction with automation", IEEE Transactions on Systems, Man, and Cybernetics, 30(3), pp. 286–297. https://doi.org/10.1109/3468.844354

Wrobel, K., Montewka, J. and Kujala, P. (2017) "Towards the assessment of potential impact of unmanned vessels on maritime transportation safety", Reliability Engineering and System Safety, 165, pp. 155–169. https://doi.org/10.1016/j.ress.2017.03.029