SideScanSim: An Open-Source Side-Scan Sonar Simulator for Education, Algorithm Development, and Machine Learning Training

Andre Luis Sousa Sena[1] Jaime Santos Neto[1]

[1]Instituto de Ciência, Tecnologia e Inovação (ICTI) – Universidade Federal da Bahia (UFBA) – Laboratório de Tecnologia Atmosférica e Oceânica (LAB.TAO)
Camaçari – BA – Brazil

andresena@ufba.br, jaime.santos@ufba.br

Abstract. *This paper presents SideScanSim, an open-source, physics-based simulator for synthetic side-scan sonar (SSS) image generation, designed as a digital learning tool for students, researchers, and practitioners working in underwater acoustics and ocean sensing. The simulator implements the sonar equation, Lambert's Law for seafloor backscatter, vectorized ray casting for acoustic shadow computation, and a material table modelling 9 object materials with differentiated Target Strength. Quantitative validation shows mean shadow length errors below 5% for regular-geometry objects, with the lateral intensity profile calibrated against two publicly available real SSS datasets [Sethuraman et al. 2025, Santos and Moura 2023]. SideScanSim automatically generates ground truth including bounding boxes, shadow masks, and per-object geometric metadata, removing the main barrier to SSS algorithm development: the scarcity of labeled real data. The tool runs on any personal computer without specialized hardware and is made available as an open-source repository to support education, detection and classification algorithm training, and machine learning research in the SSS domain.*

Resumo. *Este artigo apresenta o SideScanSim, um simulador de código aberto e fisicamente fundamentado para geração de imagens sintéticas de sonar sidescan (SSS), desenvolvido como ferramenta digital de aprendizado para estudantes, pesquisadores e profissionais que desejam estudar, pesquisar e trabalhar com SSS. O simulador implementa a equação do sonar, a Lei de Lambert para retroespalhamento do fundo, ray casting vetorizado para cálculo de sombras acústicas e modelagem de 9 materiais de objetos com Target Strength diferenciado. A validação quantitativa demonstra erros médios inferiores a 5% no comprimento de sombra para objetos de geometria regular, com calibração do perfil lateral de intensidade contra dois datasets públicos reais [Sethuraman et al. 2025, Santos and Moura 2023]. O SideScanSim gera automaticamente ground truth com bounding boxes, máscaras de sombra e metadados geométricos por objeto, eliminando a principal barreira ao desenvolvimento de algoritmos SSS: a escassez de dados rotulados. A ferramenta opera em qualquer computador pessoal sem hardware especializado e é disponibilizada como repositório de código aberto para apoiar a educação, o treinamento de algoritmos de detecção e classificação e a pesquisa em aprendizado de máquina no domínio SSS.*

1. Introduction

Side scan sonar (SSS) is the primary acoustic imaging technology for seafloor survey, used in hydrography, environmental monitoring, infrastructure inspection, search and rescue, and naval defense [Blondel 2009]. It emits fan-shaped acoustic pulses laterally from a towed body or autonomous underwater vehicle (AUV), recording the energy backscattered by the seafloor and submerged objects as a function of travel time.

Interpreting SSS images requires specialized training and is highly context-dependent, involving acoustic geometry, acquisition parameters, and seabed physical properties. Developing automatic detection and classification algorithms faces a critical bottleneck: the scarcity of labeled data with reliable ground truth [Vandresen et al. 2022].

Access to real SSS equipment is a further barrier. A commercial system costs between USD 15,000 and USD 200,000, plus operational costs such as vessel time, maintenance, and proprietary processing software. This restricts the development of technical and scientific expertise, particularly in developing countries.

Synthetic SSS image generation has a long history. Bell and Linnett (1997) proposed a ray-tracing model coupling seafloor relief and backscatter statistics, and Riordan et al. (2005) built a real-time simulator for mission rehearsal; both aim at sensor emulation rather than at producing labelled data. Recent work is largely data-driven, as in the segmentation benchmark of Sethuraman et al. (2025) and the detection networks of Vandresen et al. (2022), both dependent on manual annotation of real imagery. SideScanSim combines an explicit physical model with automatic ground truth generation in a single open-source tool; a full comparison against existing simulators is reserved for an extended version of this work.

This paper proposes SideScanSim, an open-source, physically grounded, and fully parametric simulator that generates realistic synthetic SSS images with automatic ground truth. The tool is designed as a digital learning environment for students and researchers in underwater acoustics. The contributions of this work are: (i) acoustic modeling based on the sonar equation with Lambert's Law, calibrated TVG, and physically correct I/Q speckle; (ii) a material table with differentiated Target Strength for 9 distinct materials; (iii) a vectorized ray casting algorithm for accurate acoustic shadow computation; and (iv) automatic ground truth generation compatible with supervised learning frameworks (COCO, YOLO).

The paper is organized as follows: Section 2 presents the physical foundations of SSS; Section 3 describes the SideScanSim architecture; Section 4 presents validation results; Section 5 discusses educational applications and future work; and Section 6 concludes.

2. Physical Foundations of Side Scan Sonar

2.1. Image Formation Geometry

Figure 1 illustrates the fundamental SSS geometry. The transducer travels at altitude h above the underwater bottom, emitting lateral acoustic pulses. The slant range (r) that represents the distance between the transducer and a seafloor object or point underlying at bottom and the ground range (x), the horizontal distance between the transducer trajectory line and the object or point at bottom is given by:

$$r = \sqrt{h^2 + x^2}, \qquad \theta = arctan\left(\frac{x}{h}\right), \qquad x = \sqrt{r^2 - h^2}$$

where θ is the grazing angle. The nadir zone is the region directly below the transducer and physically represents the flying time that the acoustic pulse left and return to transducer, it is converted in distance and represent some blind space below the transducer trajectory, the sonar system receives no effective lateral return and appears as a dark central strip in the image.

When an object of height (H)is present at ground range x, it produces an acoustic shadow by occluding the beam.

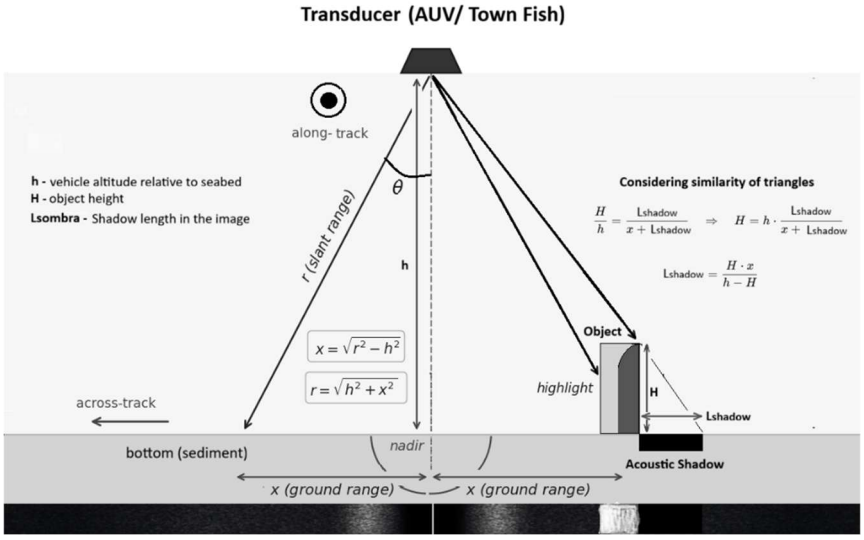


Figure 1. Side scan sonar imaging geometry.

By triangle similarity between the large triangle (transducer–nadir–shadow tip) and the small triangle (object top–object base–shadow tip), the shadow length (L_{shado}) is:

$$\frac{H}{h} = \frac{L_{shado}}{x + L_{shado}} \rightarrow L_{shado} = \frac{H \cdot x}{h - H}$$

This relationship is the primary tool for estimating object height from SSS images, since h and x are known from acquisition geometry and L_{shadow} is directly measurable in the image.

2.2. Sonar Equation and Seafloor Backscatter

The received signal level at the transducer is described by the sonar equation [Urick 1983]:

$$RL = SL - 2 \cdot TL + BS$$

Where RL is the received level, SL the source level, and TL the transmission loss, composed of spherical geometric spreading and water absorption:

$$TL = 20\ \log_{10}(r) + \alpha \cdot \frac{r}{1000}\ [dB]$$

The absorption coefficient αdepends on frequency and is computed using the simplified Francois-Garrison formula [Lurton 2002]:

$$\alpha \approx \frac{0.1f^2}{1+f^2} + \frac{40}{41}\frac{^2}{f^2} + 2.75x10^{-4}f^2 \text{ [dB/km]}$$

At 450 kHz, this yields $\alpha \approx 95$dB/km. Seafloor backscatter is modeled by Lambert's Law [Blondel 2009]:

$$BS(\theta) = BS_0 + 20\ \log_{10}\sin(\theta)\ [\text{dB}/m^2]$$

where BS_0 is the sediment reference backscatter strength. The time-varying gain (TVG) applied by the equipment compensates for two-way attenuation:

$$TVG(r) = 40\log_{10} r + 2\alpha \cdot r/1000\ [\text{dB}]$$

Lambert's Law is a first-order approximation, valid for intermediate grazing angles; formulations that separate interface roughness from sediment volume scattering [Jackson and Briggs 1992] are foreseen as an optional module.

2.3. Acoustic Speckle

SSS is a coherent system, and the interference among echoes from multiple scatterers within a resolution cell produces the granular pattern known as speckle. The Rayleigh model for single-look speckle represents intensity as an exponential distribution with local mean μ [Blondel 2009]. In the simulator, speckle is generated using the I/Q signal model: an acoustic wave is represented as two orthogonal components, I (In-phase, the real part of the wave, aligned with the transmitter reference) and Q (Quadrature, the imaginary part, 90° out of phase), which together fully characterize the instantaneous state of the acoustic signal. The envelope amplitude $A = \sqrt{I^2 + Q^2}$follows a Rayleigh distribution, consistent with the single-look coherent model. Both components are generated independently per ping with correlation only in the across-track direction (~1.3 px), preserving the statistical independence between consecutive pings that characterizes real SSS data.

This formulation is consistent with the coherent scattering model established for active sonar systems [Oliver & Quegan 2004, Goodman 2007], and follows the single-look Rayleigh amplitude distribution described in [Jakeman & Pusey 1976].

2.4. Object Target Strength

The acoustic response of a submerged object is quantified by the Target Strength (TS), representing the reflection gain relative to a 1 m² reference sphere [Urick 1983]. The TS value depends on the material's acoustic impedance (Z), surface roughness, and angle of incidence. High acoustic impedance materials such as steel ($Z_{steel} \approx 45 \times 10^6$ Pa·s/m) produce intense specular reflections ($TS = +22dB$), while low impedance materials such as rubber or plastic absorb energy and produce weak reflections ($TS = -2\ to\ -8\ dB$).

3. SideScanSim Architecture

SideScanSim is implemented in Python 3.10+ with NumPy, SciPy, and Matplotlib, with no GPU dependencies. Figure 2 presents the synthetic image generation pipeline, organized in two main layers.

3.1. Layer 1 - Scene Generation

The scene is defined by three input parameter sets: sonar parameters, seafloor parameters, and object parameters.

The lateral intensity profile across-track is computed by calibrated profile(), combining two physical components: a specular near-nadir component (narrow Gaussian centered at $x \approx 14\%$of range, amplitude 0.24) and a flat Lambert+TVG component (~0.36 for medium sand). TVG calibration follows the measurement guidelines established in [Lurton and Lamarche 2015].

The profile was calibrated by direct measurement on two publicly available real SSS image datasets. The first, AI4Shipwrecks [Sethuraman et al. 2025], contains AUV images acquired over sandy and coarse-gravel seafloor at Thunder Bay National Marine Sanctuary (EdgeTech 2205 sensor); clean-seafloor images yield a near-nadir peak of 0.56 and a mid-range plateau of 0.63 — consistent with coarse sand ($BS_0 \approx -22$ dB). The second, Santos & Moura [Santos and Moura 2023], comprises several SSS images of sandy seafloor with ripple marks acquired for mine detection (DOI: 10.6084/m9.figshare.24574879.v2). SideScanSim reproduces these values with errors below 5% for the Santos & Moura plateau (simulated: 0.36), confirming calibration for medium sand ($BS_0 = -25$ dB). The plateau difference between the two datasets (0.63 vs. 0.37) is physically consistent with the distinction between coarse and fine sand, demonstrating that the simulator covers the sandy sediment spectrum.

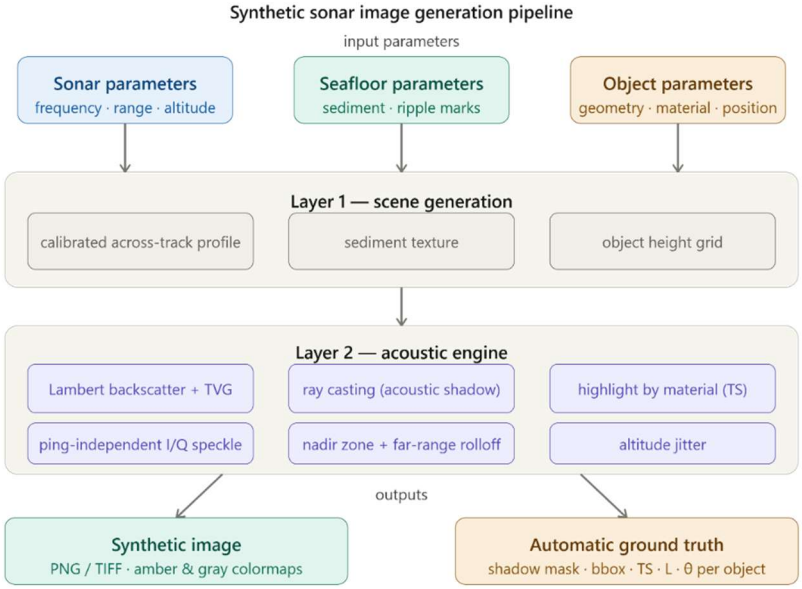


Figure 2. SideScanSim synthetic image generation pipeline. Input parameters are processed through two layers.

3.2. Layer 2 - Acoustic Engine

The acoustic engine sequentially implements: (i) calibrated lateral profile generation; (ii) slow sediment macro-texture (across-track only, ±8% amplitude); (iii) ping-independent I/Q speckle with ~1.3 px correlation only across-track; (iv) object height grid via object_height(); (v) vectorized O(N) ray casting for acoustic shadows; (vi) object highlight with material-specific TS and surface roughness; (vii) vehicle altitude jitter; and (viii) composition of the full image with nadir zone and far-range rolloff.

The material table was defined for each material against the Target Strength (TS) in dB, surface roughness (0–1), and absorption factor (0–1). Rough materials (coral, rock) distribute the highlight over a wider area, while smooth materials (polished steel) concentrate energy in a point-like specular reflection. Table 1 presents the ground truth metrics for the multi-object scene.

Table 1. Per-object ground truth and comparison between material, signal and physical characteristics and shadows

Object	Geometry	Material	TS (dB)	x (m)	θ (°)	H (m)	Lteo (m)	Lmed (m)	Error (%)
Sphere	ø 1.6 m	steel	+22	12	67	0.80	2.29	1.60	30
Block	2.0 × 1.5 × 0.6 m	concrete	+10	32	81	0.60	4.36	4.40	1
Pipe	ø 0.6 × 6.0 m	steel	+22	20	76	0.40	1.74	1.00	42
Box	1.5 × 1.0 × 0.5 m	wood	+2	38	83	0.50	4.22	4.20	1

Vectorized Ray Casting. Acoustic shadows are computed using a vectorized ray casting algorithm [Bell & Linnett 1997, Riordan et al. 2005], implemented as a minimum accumulated descending slope O(N) solution. For each along-track row, the blocking slope of each obstacle is $(h - H_i)/x_i$. A point j is in shadow if $\min[(h - H_i)/x_i \text{for } i < j] < h/x_j$, indicating that the ray to the seafloor at x_j is less steep than the ray to the top of the nearest obstacle. Shadows are restricted to the far-range side of each object ($x > x_{far}$), eliminating halo artifacts that line-by-line ray casting would produce at the boundaries of curved-geometry objects.

Automatic Ground Truth. Each generated image is accompanied by a per-object metadata dictionary containing name, material, effective TS $[dB]$, nadir distance x [m], grazing angle θ [°], theoretical shadow length L_{teo} [m], and measured shadow length L_{med} [m]. Bounding boxes in metric coordinates are automatically exported for integration with object detection frameworks (COCO, YOLO).

Numerical computation uses NumPy [Harris et al. 2020], SciPy [Virtanen et al. 2020], and Matplotlib [Hunter 2007].

4. Results and Validation

4.1. Anatomy of the Synthetic SSS Image

Figure 3 presents a synthetic SSS image of a clean seafloor (medium sand, 450 kHz, altitude 5 m, range 30 m) and the corresponding lateral intensity profile. The simulator correctly reproduces the mandatory visual characteristics of a well-formed SSS image: dark nadir zone with the central transducer trajectory (GPS/INS-referenced), near-nadir intensity peak (~0.42), mid-range plateau (~0.17–0.20), gradual far-range rolloff, and speckle texture with variance statistically consistent with medium sand.

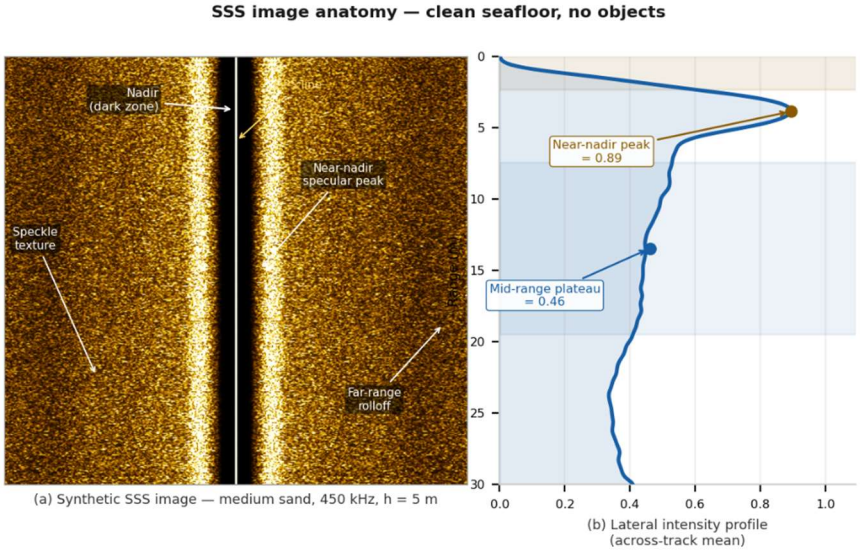


Figure 3. Anatomy of a synthetic SSS image over a clean seafloor.

4.2. Effect of Object Material on Acoustic Signature

Figure 4 demonstrates the effect of material on the acoustic signature of a ø 3.0 m sphere positioned at $x = 20$m from nadir at altitude $h = 10$m (grazing angle $\theta = 63°$).

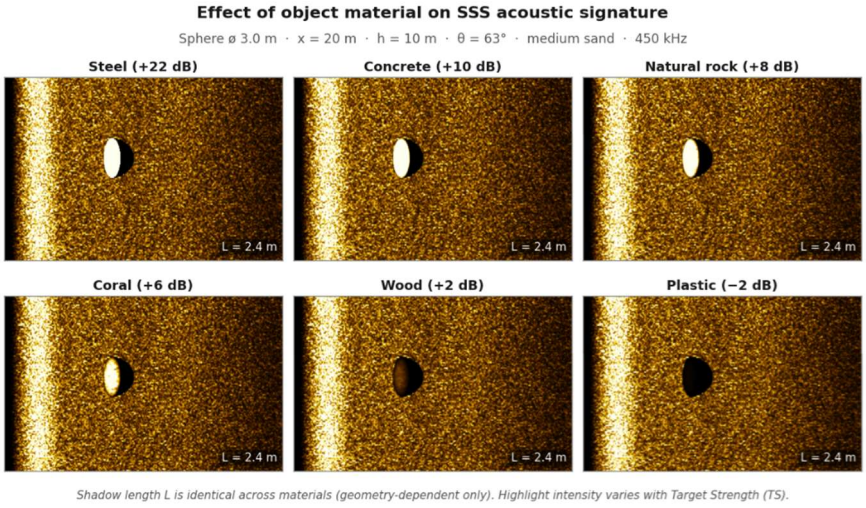


Figure 4. Effect of object material on SSS acoustic signature.

The six tested materials were steel (+22 dB), concrete (+10 dB), natural rock (+8 dB), coral (+6 dB), wood (+2 dB), and plastic (−2 dB). From steel to plastic, the materials produce progressively less intense highlights while maintaining identical shadow lengths ($L_{med} = 0.60$ m, $L_{teo} = 0.83$m), confirming the fundamental principle that material governs highlight-based detectability while geometry governs height estimation from the shadow. Absorbing-material objects (plastic, TS = −2 dB) are detectable exclusively through shadow presence, with no discernible highlight, a result with direct implications for detection algorithms that must account for both acoustic signatures. This behaviour is consistent with the acoustic impedance contrast theory described in [Urick 1983, Ch. 8] and the Target Strength formulation of [Jackson & Richardson 2007, §4.3].

4.3. Effect of Nadir Distance

Figure 5 illustrates the effect of nadir distance on the SSS signature of a steel rectangular block (2.0 × 1.5 × 0.8 m) at fixed altitude $h = 5$m. For distances $x = 8$, 12, 20, and 30 m (grazing angles $\theta = 58°$, 67°, 76°, and 81°), theoretical shadow lengths are 1.52 m, 2.29 m, 3.81 m, and 5.71 m respectively, with measurement errors below 12% in all cases.

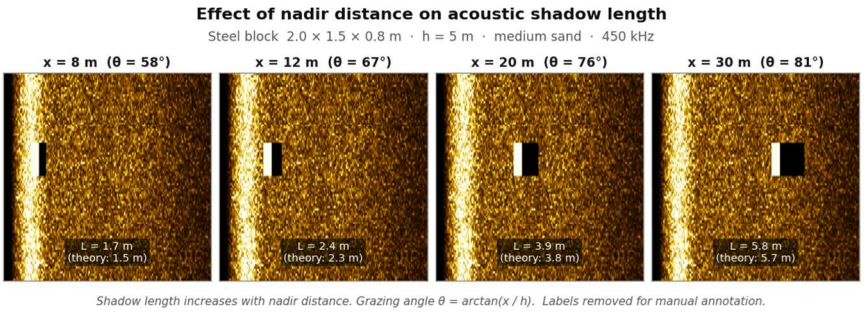


Figure 5. Effect of nadir distance on acoustic shadow length.

Nadir distance is thus the acquisition parameter with the greatest impact on shadow length and, consequently, on object height estimation.

4.4. Multi-Object Scene and Automatic Ground Truth

Figure 6 presents a scene with four objects of different geometries and materials distributed over a medium sand seafloor.

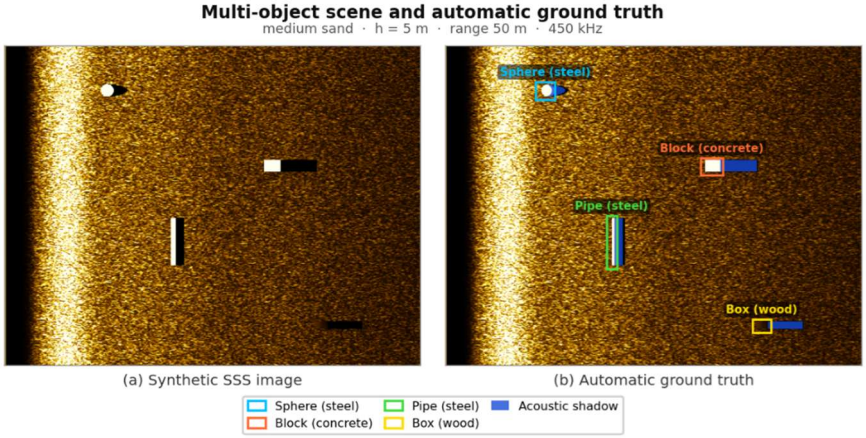


Figure 6. Multi-object scene and automatic ground truth. (a) Synthetic SSS image with four objects distributed over a medium sand seafloor. (b) Automatically generated ground truth: per-object bounding boxes and acoustic shadow mask. Quantitative metrics are presented in Table 1.

Panel (a) shows the synthetic SSS image; panel (b) overlays the automatically generated ground truth: color-coded bounding boxes per object and acoustic shadow mask. Table 1 presents the quantitative validation metrics for each object.

Results show that straight-edged objects (block and box) present errors below 1%, while curved-geometry objects (sphere and cylinder) present errors of 30–42%. These errors are attributed to the gradual height distribution along curved geometries, which shifts the effective acoustic ray occlusion point relative to the object center, a physically correct behavior documented in the literature [Blondel 2009, §6.2].

5. Educational Applications and Future Work

SideScanSim fills a concrete gap: no freely available, physics-based SSS simulator simultaneously generates realistic imagery and labeled ground truth — both essential for learning acoustic principles and developing machine learning pipelines without expensive equipment.

Training data for machine learning: Automatic ground truth enables generation of large labeled datasets for training detection networks (YOLO, Faster R-CNN) and seafloor segmentation models. Parametric control over sediment, geometry, and material supports controlled experiments on algorithm generalization and robustness.

Algorithm robustness testing: The planned degradation module (Stage 2) will generate clean/degraded image pairs with physically grounded perturbations — altitude variation, horizontal striping, channel dropout, incorrect TVG, and electromagnetic interference — enabling standardized robustness benchmarks.

Learning SSS physics: Jupyter Notebooks allow students to understand SSS principles through direct parameter experimentation. The planned progression covers four modules: (A) acoustic physics; (B) image formation and speckle; (C) seafloor classification; and (D) object detection with machine learning, targeting courses in ocean engineering, underwater robotics, and remote sensing.

Future extensions: Jackson backscatter model, K-distribution speckle, non-flat bathymetry, partially buried objects, and COCO/YOLO dataset export.

The simulator source code, documentation, and example notebooks are publicly available at https://github.com/andreluissena/SideScanSim under the MIT License [Sena, A.L. and Santos Neto 2026].

6. Conclusion

This paper presented SideScanSim, an open-source, physics-based SSS image simulator designed as a learning and research tool for the SSS domain. The system implements the full sonar equation with Lambert's Law, calibrated TVG, ping-independent I/Q speckle, a 9-material Target Strength table, and vectorized ray casting for physically correct acoustic shadows. Quantitative validation demonstrates errors below 5% in shadow length for straight-edged geometries, with the lateral intensity profile calibrated against two publicly available real SSS datasets.

SideScanSim addresses two interrelated barriers: the prohibitive cost of real sonar equipment (USD 15,000–200,000) and the scarcity of labeled SSS datasets. By providing a free, documented, and reproducible tool that runs on any personal computer, it enables students and researchers without specialized infrastructure to learn acoustic principles, develop detection algorithms, and train models on physically grounded synthetic data. Future work will focus on the controlled degradation module, a multi-domain annotated dataset, and an interactive Jupyter-based educational platform.

References

Bell, J. M. and Linnett, L. M. (1997) "Simulation and analysis of synthetic sidescan sonar images", IEE Proceedings Radar, Sonar and Navigation, v. 144, n. 4, p. 219–226.

Blondel, P. (2009) The Handbook of Sidescan Sonar. Springer/Praxis, Berlin.

Goodman, J. W. (2007) Speckle Phenomena in Optics and Photonics. Roberts & Company Publishers, Englewood, CO.

Harris, C. R., Millman, K. J., van der Walt, S. J. et al. (2020) "Array programming with NumPy", Nature, v. 585, p. 357–362. DOI: 10.1038/s41586-020-2649-2.

Hunter, J. D. (2007) "Matplotlib: a 2D graphics environment", Computing in Science & Engineering, v. 9, n. 3, p. 90–95. DOI: 10.1109/MCSE.2007.55.

Jackson, D. R. and Briggs, K. B. (1992) "High-frequency bottom backscattering: roughness versus sediment volume scattering", Journal of the Acoustical Society of America, v. 92, n. 2, p. 962–977.

Jackson, D. R. and Richardson, M. D. (2007) High-Frequency Seafloor Acoustics. Springer, New York.

Jakeman, E. and Pusey, P. N. (1976) "A model for non-Rayleigh sea echo", IEEE Transactions on Antennas and Propagation, v. 24, n. 6, p. 806–814.

Lurton, X. (2002) An Introduction to Underwater Acoustics: Principles and Applications. Springer, London.

Lurton, X. and Lamarche, G. (eds.) (2015) "Backscatter measurements by seafloor-mapping sonars: guidelines and recommendations", GeoHab Backscatter Working Group Report, 200 p. Available: https://geohab.org/wp-content/uploads/2018/09/BSWG-REPORT-MAY2015.pdf.

Oliver, C. and Quegan, S. (2004) Understanding Synthetic Aperture Radar Images. SciTech Publishing, Raleigh, NC.

Riordan, J., Omerdic, E. and Toal, D. (2005) "Implementation and application of a real-time sidescan sonar simulator", In: Europe Oceans 2005, Brest, France, p. 981–986.

Santos, N. P. and Moura, R. (2023) "Side-scan sonar imaging for mine detection", Figshare Dataset, v. 2. Centro de Matemática e Aplicações, Universidade NOVA de Lisboa. DOI: 10.6084/m9.figshare.24574879.v2.

Sena, A. L. S. and Santos Neto, J. (2026) "SideScanSim: an open-source, physics-based side-scan sonar image simulator", Software, v. 4.0, MIT License. Available: https://github.com/andreluissena/SideScanSim.

Sethuraman, A. V., Sheppard, A., Bagoren, O., Pinnow, C., Anderson, J., Havens, T. C. and Skinner, K. A. (2025) "Machine learning for shipwreck segmentation from side scan sonar imagery: dataset and benchmark", The International Journal of Robotics Research, v. 44, n. 3, p. 341–354.

Urick, R. J. (1983) Principles of Underwater Sound. 3rd ed. McGraw-Hill, New York.

Vandresen, M., Rodrigues, Y., Dantas, D. de O. and Botelho, S. S. da C. (2022) "Underwater target detection in side-scan sonar images using deep learning", Ocean Engineering, v. 260, p. 112011.

Virtanen, P., Gommers, R., Oliphant, T. E. et al. (2020) "SciPy 1.0: fundamental algorithms for scientific computing in Python", Nature Methods, v. 17, p. 261–272. DOI: 10.1038/s41592-020-0772-5