



European
Commission

Horizon Europe
Work Programme 2021-2022
Digital, Industry and Space



A **M**ULTIPLEXED PLASMO-PHOTONIC **B**IOSENSING
PLATFO**R**M FO**R** RAPID AND INTELLIGENT **S**EPSIS
DI**A**GNOSIS AT THE POINT-OF-CARE

D2.3 Final report on sepsis detection and classification system/component specifications and designs

Project Information

Project name:	A M ultiplexed plasmo-photonic B iosensing platfo R m fo R rapid and intelligent S epsis d I A gnosis at the point of care
Project acronym:	AMBROSIA
Project start date:	1/1/2023
Project duration:	48 months
Contract number:	101093166
Project coordinator:	Nikolaos Pleros / AUTH
Topic:	HORIZON-CL4-2022-DIGITAL-EMERGING-01-03: Advanced multi-sensing systems

Document Information

Document number and title:	D2.3 Final report on sepsis detection and classification system/component specifications and designs.
Document type:	R-Document, report
Contractual date of delivery:	30/06/2025
Calendar date of delivery:	04/08/2025
Work package number:	WP2 System architecture and subsystem conceptual design and specifications
Lead partner:	AUTH
Dissemination level:	PU
Document status:	Version 1, final

Author List:

Organization	Author
AUTH	Evangelia Chatzianagnostou, Stylianos (Stelios) Simos, Konstantinos Fotiadis, Dimosthenis Spasopoulos, Lamprini Damakoudi, Anastasios Tefas, Dimitris Spanos, Nikolaos Passalis, Panagiotis Zdoupas
ICN2	M.-Carmen Estevez
VHIR	Juan José González López, Juan Carlos Ruiz Rodríguez
MILQ	Ikerne Etxebarria
ARGO	Michał Szaj
SMART	Mack Johnson, Genia Patent
XCEL	Alex Farrell, Alin Fecioru, Ruggero Loi
BIOPIX-T	Stella Chatziioanidou, George Papadakis
SOTON	T. Dominguez Bucio, I. Chakraborty, X. Yu, I. Johnson, A. Shoa, Q. Luo, T. Rangarajan, E. Sandell, I. Zeimpekis, F. Gardes
LGT	Jose Carreira
UB	Jean-Claude Weeber
UOI	Dimitris Bellas, Eleftheria Lampadariou, Elefterios Lidorikis

TABLE OF CONTENTS

1. EXECUTIVE SUMMARY	4
2. INTRODUCTION	5
2.1. PURPOSE OF THIS DOCUMENT	5
2.2. DOCUMENT STRUCTURE	5
2.3. AUDIENCE	5
3. TECHNICAL SPECIFICATIONS AND DESIGNS OF THE AMBROSIA SUBSYSTEMS 6	
3.1. PLASMO-PHOTONIC SENSORS	6
3.2. MICROFLUIDICS	12
3.3. INP ACTIVE ELEMENTS AND INTERFACE WITH THE SiN PLATFORM	14
3.3.1. InP active elements	14
3.3.2. InP-to-SiN interface	17
3.4. PHOTONIC NEUROMORPHIC CIRCUIT	19
3.5. ELECTRONICS	21
3.6. READ-OUT SYSTEM	24
3.7. PoC PLATFORM	25
4. USE CASES AND ARTIFICIAL INTELLIGENCE TRAINING MODELS FOR SEPSIS DETECTION AND CLASSIFICATION	27
4.1. SEPSIS SCENARIOS AND TARGETED SET OF BIOMARKERS AND PATHOGENS	27
4.2. SEPSIS DIAGNOSIS AND CLASSIFICATION DATA	29
4.3. AI TRAINING MODELS	31
5. CONCLUSIONS	38

1. Executive summary

This deliverable reports on the final designs and technical specifications of the subsystems and components comprising the AMBROSIA biosensing platform. Minor readjustments of some sub-modules' specifications (like the InP actives) may be needed according to the feedback received from the experimental results, while the assembly process flow of the final prototype will be updated accordingly. Clinical aspects, such as the set of the targeted sepsis-related biomarkers and pathogens, as well as the assay ranges requirements for different clinical scenarios, are also included. Finally, the deliverable describes the developed artificial intelligence (AI) training models to be used for sepsis detection and severity classification utilizing photonic neural networks.

2. Introduction

2.1. Purpose of this document

The objective of this deliverable is to report the final technical specifications of the subsystems and components of the AMBROSIA biosensing platform including the AI training models.

2.2. Document structure

The present deliverable is split into the following major parts:

- Technical specifications and designs of the AMBROSIA subsystems
- Use cases and AI training models for sepsis diagnosis and classification
- Conclusions

2.3. Audience

The content of this document is public.

3. Technical specifications and designs of the AMBROSIA subsystems

3.1. Plasmo-photonic sensors

✚ Sensors without Bragg-gratings

As previously reported, two sensor configurations have been investigated, both numerically and experimentally: (i) a SiN-based, **power-balanced plasmo-photonic Mach-Zehnder Interferometer (MZI)** with aluminum (Al) plasmonic stripe waveguides on both arms, and (ii) a **single-arm bimodal plasmo-photonic interferometer**. Both configurations were first evaluated with bare plasmonic stripes. Afterwards, dielectric Bragg-gratings were considered on top of the metallic stripes in the MZI structure to further enhance sensitivity.

The conceptual layouts of both configurations are depicted in Figure 1 and Figure 2.




Figure 1: (a) Conceptual 3D schematic of the plasmo-photonic MZI sensor, Cross section of (b) the input photonic WG, (b) the photonic WG after the linear taper and (c) the aluminum plasmonic WG.




Figure 2: (a) Conceptual 3D- schematic of the bimodal plasmo-photonic sensor, cross section of (b) the Al waveguide with slab-based SiO₂ bottom layer.

The optimal cross-section of the **bimodal interferometer** comprises two access, 800nm-thick SiN waveguides separated by a thin aluminum (Al) metal stripe located on top of a 320nm-thick, infinite SiO₂ layer (bottom layer) forming two metal-insulator interfaces that are both able to support plasmonic modes (top and bottom). The propagation losses were calculated equal to $0.07 \text{ dB}/\mu\text{m}$ for the top plasmonic mode and $0.064 \text{ dB}/\mu\text{m}$ for the bottom plasmonic mode while the losses at the interface between the photonic and plasmonic parts were $\sim 5 \text{ dB/per interface}$. The lengths chosen for fabrication were between 110 and 200 μm and were selected based on the corresponding free spectral range (FSR) and total insertion losses. For this configuration, the **maximum simulated bulk refractive index (RI)**

sensitivity was equal to 13123 nm/RIU irrespective of the sensing length corresponding to an expected limit of detection of 7.6×10^{-7} , when considering a system resolution of 10 pm.

The characterization at AUTH showed that all fabricated bimodal sensors, where the metal was deposited either at LGT or UB¹, are functional and show a very good agreement between simulated and experimental characteristics, and very good repeatability between different sensors with the same structural parameters. The best performance was observed for sensors with a plasmonic length of 160 and 190 μm and the **maximum experimentally achieved sensitivity was 13312 nm/RIU**.

Regarding the **MZI configuration**, the optimal cross-section of the plasmonic part comprises a cavity down to the SiN and a part of the bottom oxide layer on which the Al plasmonic stripe is deposited to introduce a vertical offset (VO) between the two waveguides (photonic and plasmonic). Using 3D FDTD simulations, the optimum value of the VO was found equal to 350 nm resulting in 40% coupling efficiency or interface losses of 4 dB/interface. The propagation losses of the plasmonic mode were calculated equal to 0.06 dB/ μm . Each branch of the MZI contains heater-based phase shifters in order to tune the wavelength dip within the desired spectral window (C-band). Plasmonic stripes of 70 μm and 140 μm length have been selected for fabrication. The specific variants that have been included in the lithography mask were selected to experimentally demonstrate (i) the dependence of sensitivity on the sensor FSR, and (ii) the dependence of sensitivity on the sensing area length. Thus, 4 sensor variants with the same 70 μm long plasmonic sensing area and varying length difference (ΔL) between the two branches leading to varying FSRs (70 nm, 100 nm (x2), 186 nm, and >300 nm) as well as 2 sensor variants with the same FSR of 186 nm but different plasmonic sensing area length (70 μm and 140 μm) have been included. The **maximum bulk RI sensitivity was calculated equal to 28000 nm/RIU** for the sensor with >300 nm FSR by means of numerical circuit-level simulations leading to an expected limit of detection of 3.5×10^{-7} .

The characterization at AUTH showed that all fabricated MZI sensors are also functional and show a very good agreement between simulated and experimental characteristics. The best performance was observed for sensors with a length difference of 5 μm and the **maximum experimentally achieved sensitivity was 20605 nm/RIU**. Finally, experimental characterization of the phase-shifters in all MZIs showed that a phase shift of 2π can be achieved in all cases. The tuning efficiency for an MZI with ~200 nm FSR was measured equal to 0.3 nm/mW when a voltage is applied to the phase shifter of the upper branch, and 0.69 nm/mW when a voltage is applied to on the phase shifter of the bottom branch which is closer to the contact pads.

Table 1 summarizes the technical specifications and the theoretical performance metrics of both configurations in comparison to the experimental ones.

Table 1. Technical specifications and performance metrics (both theoretical and experimental) of the AMBROSIA sensors.

Parameter	Balanced MZI		Bimodal	
	Simulation	Experiment	Simulation	Experiment
Photonic WG dimensions	1 μm $\times$ 800 nm		1 μm $\times$ 800 nm	

¹ In the **first fabrication run**, the fabrication was split into **4 sublots**: initially one for the bimodal and one for the MZI configuration since the two different sensor configurations have different cross-sections, and, soon after, the sublots were further split into plasmonic metal done at LGT and at UB to derisk a potential bottleneck. Additionally, two different cavity configurations have been investigated at LGT: (i) one with vertical sidewalls which is closer to the simulated one, but leads to metal residues on the sidewall, and (ii) one with sloped sidewalls that enables better photoresist coverage and, thus, prevents metal deposition on the sidewall. For the sublots where the plasmonic metal is defined at LGT, **it has been decided to proceed with the sloped cavity profile** after validation of proper performance of a look-ahead wafer at AUTH (more details can be found in D3.1 and D3.2).

$(W \times t)$				
Plasmonic WG dimensions $(W \times t)$	7 $\mu\text{m} \times 80\text{ nm}$		7 $\mu\text{m} \times 80\text{ nm}$	
Vertical Offset (nm)	350 nm		-	
Interface losses (dB)	4/interface	-	5/interface	-
Plasmonic propagation losses (dB/ μm)	In water 0.06	-	In water 0.07 (top) 0.064 (bottom)	In air 0.057 (top) 0.054 (bottom)
ΔL (μm)	2 – 22	5	-	
Bottom layer thickness (nm)	-		320 nm	
Plasmonic WG length (μm)	70 / 140	70	100 – 200	160, 190
FSR (nm)	>300 – 47	200	199 – 100	>100
Tuning efficiency of the phase shifter	-	2pi	-	-
ER (dB)	50 – 60	30	29 – 64	25
Sensitivity (nm/RIU)	2067 – 28000	20605	13123	13312

MZI sensors with Bragg-gratings

To enhance the sensitivity of the plasmo-photonic sensors, the UOI has undertaken the design and simulation of Bragg grating configurations on the metal waveguide. These structures introduce defect modes that enable the exploitation of the slow light effect, significantly improving sensor performance. The adjustable structural parameters, which were optimized, include the number of Bragg periods, the cavity length d_{cav} , and the tooth height, as illustrated in Figure 3. Two materials were investigated for the Bragg gratings: silicon nitride (Si_3N_4) and parylene. The final design specifications for both materials are summarized in Figure 3b. The simulation results indicate that the scattering losses are strongly influenced by the tooth height. In particular, increasing the tooth height leads to reduced scattering, thereby preserving high sensitivity due to the enhanced performance of the Bragg gratings. Nevertheless, even at finite tooth heights, the sensor sensitivity remains significantly higher than that of a bare plasmo-photonic sensor, as detailed in Figure 3c. The complete optimization procedure and final design specifications are thoroughly documented in Deliverable D3.1.




Figure 3: Design and optimization of Bragg gratings for the MZI sensor. (a) Illustration of the Bragg cavity integrated on the plasmonic slab layer. (b) Specifications of the optimized Bragg grating designs. (c) Simulated sensitivity of the Si_3N_4 -based Bragg grating.

The fabrication and comprehensive optical characterizations of SiN-based Distributed Bragg Reflectors (DBRs) are reported in detail in D3.2 by UB. Based on experimental results, it is clear that the performance of deep-etched SiN DBRs for surface plasmon polariton can be significantly affected by scattering process on each ridge comprising the DBR. For optimized performances, it is highly desirable to minimize the so-called parasitic scattering process of the incident SPP. Figure 4(A) shows the stack of material used to fabricate deep-etched SiN DBR. Due to the specific fabrication process (see D3.2), a narrow gap remains at the end of the fabrication process between the upper surface of the plasmonic metal and the bottom surface of the SiN layer. This gap plays a critical role in the aforementioned parasitic scattering process, as it determines to a large extent the profiles of the modes supported by the metal/Low Index/High Index (M/Li/Hi) stack.

Figure 4(B) and (C) show the modes profiles of the fundamental TM modes supported by the M/Li/Hi stack along with the mode profile of the incident Water/Au interface SPP in case of a water gap with a thickness of 50 nm and 11 nm respectively. It is apparent from the mode profiles that if the water gap is




Figure 4: (A) Computational situation showing a thick-SiN layer on top of a plasmonic metal. A gap filled with water is created between the metal and SiN due to the current fabrication process. (B) The water gap is 50 nm thick. Mode profile matching between the incident SPP and the fundamental TM0 is very poor. (C) For a water gap of 11 nm the mode profile matching is much more convenient. The same result is observed by replacing the 11nm-water gap by a 15nm-SiON gap.

50nm-thick (expected to correspond to the current experimental situation), the two modes of interest (incident SPP and fundamental TM0 mode of the M/Li/Hi stack) feature a very different mode profile leading to parasitic scattering. However, if the gap is reduced down to 11 nm, a much more convenient mode profile matching is observed suggesting that for improved performances, one should reduce the water thickness below the SiN layer down to 11 nm. Such a thickness is out of the reach of the current fabrication process when considering a water gap, however the same matching of mode profiles can be obtained when replacing the water gap by a silicon oxynitride thin layer with a refractive index of 1.46 (at 1550 nm). In this case, the optimal mode matching is observed at a thickness of 15 nm (instead of 11 nm when the low index material is water).

In order to evaluate the importance of such a low index spacer for the properties of SPP DBRs, we numerically investigate the properties of a single SiN ridge with a trapezoidal shape (slant angle 88° corresponding to the experimental situation, see D3.2) and a height of $3.5\mu\text{m}$. Figure 5 (B)-(E) show the change of SPP Reflectivity, Transmissivity and Scattering Loss as a function of the width of the SiN ridge in case of a 50nm-thick gap of water ((B) and (C)) and for a 15 nm gap of SiON ((D) and (E)). As expected from the mode profile analysis, a slight decrease of scattering losses is observed in the presence of the SiON mode profile adapter for wide SiN ridges ($1.5\mu\text{m}$ - $2.5\mu\text{m}$), however this improvement is negligible for a single ridge with a width of about $1.0\mu\text{m}$ which is the SiN ridge width targeted experimentally (see D3.2).




Figure 5: (A) Computational situation for the a single SiN ridge illuminated by an incident SPP. (B) and (C) A 50nm-water gap exists beneath the SiN layer. (D) and (E) A 15-SiON gap exists beneath the SiN layer. The mode profile matching has very little contribution on SPP scattering for a single narrow SiN ridge.

Beyond the situation of a single ridge, it is worth considering DBR as well. For a free-space wavelength of 1550 nm, and DBR with a period of 5.40 μm, a width for the ridges of 1.05 μm leads to a maximum SPP reflectivity. Figure 6 shows the change of reflectivity, Transmissivity and Losses as a function of the number of ridges comprising the DBR. As shown in Figure 6(A) and (C), the presence of a SiON spacer does not change significantly the performance of the DBR. Indeed, whatever the spacer (50 nm H₂O or 15 nm SiON), the field distributions in Figure 6(B) and (D) display neat scattering losses. Based on those results, we conclude that unfortunately, in our case, an adjustment of the mode profile between the incident SPP and the fundamental TM₀ mode of the M/Li/Hi stack does not contribute to a decrease of parasitic scattering losses. The origin of this result has to be traced back to the coupling of the incident SPP to the modes of the M/Li/Hi stack. In our case, the stack we consider is several microns high and hence supports many TM polarized modes. As a result, the incident SPP excites not only the fundamental TM₀ mode but also many other modes (TM₁, TM₂...). In order to demonstrate this multi-mode excitation, we consider in Figure 7 the situation of a single bulky SiN pad featuring a mesoscopic width. Unlike the situation of narrow SiN ridges comprising a DBR, we observe now in Figure 7(A) a clear benefit of the presence of a 15 nm SiON spacer beneath the SiN layer with scattering losses that are more than three-times smaller as compared the 50 nm water gap. The electric field distributions displayed in Figure 7(B) and (C) are obtained for a pad width of 17 μm and with a 50nm-water or a 15nm-SiON gap respectively.




Figure 6: Change of a R, T, and Loss (L) as a function of the number of ridges in a Reflecting DBR (see text). (A) and (B) A 50nm-gap of water is present under the SiN layer. (C) and (D) A 15nm-SiON mode profile adapter layer is used. The profile adjustment plays a minor role in the performances of the DBR if the ridges comprising the DBR are narrow (around $1\mu\text{m}$).




Figure 7: (A) Scattering losses for a single mesoscopic SiN pad on top a plasmonic film. (B) (resp. (C)) Electric field norm distribution for a SiN pad with a width of $17\mu\text{m}$ if a 50 nm gap of water (resp. 15nm-SiON gap) exists beneath the SiN layer.

In Figure 7(B), the multimode excitation has a dramatic impact on the field distribution at the output of the SiN pad whereas in the presence of the 15nm-SiON gap, the dominant excitation of the TM₀ fundamental mode of the SiN pad allows for an efficient re-excitation of the SPP at the output and a corresponding large decrease of scattering losses in this case.

In summary, UB has investigated numerically the benefit of the presence of a low index spacer between the plasmonic metal and the thick SiN layer with the aim of reducing parasitic scattering of the incident SPP. It turns out that, for the situation we consider, scattering losses are mostly governed by a multimode beating pattern originating for the multi-mode excitation of the SiN pad by the incident SPP.

For narrow ridges, this beating does not contribute significantly to scattering owing to the beating period of several microns and, thus, the presence of a mode profile adapter (SiON layer) is only of little practical interest for DBR. On the contrary, when bulky configurations are considered, the situation is changed provided the length of the bulky pad is long enough for the beating pattern to develop. In this case, the presence of the SiON thin layer is of key interest and is currently investigated in the context of SiN-based coplanar micro-optics for SPPs such as micro-prisms and micro-lenses.

3.2. Microfluidics

As described in D4.1, MLIQ developed the microfluidics chip for the sensor chip of the 1st fabrication run based on the specifications provided by other project partners (AUTH, LGT, ICN2, and ARGO). The sensor chip dimensions are 22 x 17 mm² (specifically 21.88 x 16.91 mm²), with 22 x 15.5 mm² for the microfluidic area. The remaining 1.5 x 22 mm² section is reserved for PCB connections.


The material of choice was PDMS. To accommodate the overall chip dimensions and maintain a minimum spacing of 1.65 mm between channels, the microfluidic chip contains 8 channels. Each channel has 14 mm in length and 250 microns in width and is equipped with dedicated inlets and outlets to allow the measurement of different indicators and to support selective functionalization. To enhance accessibility and usability, the inlets and outlets were deliberately offset, allowing for direct connection of microfluidic tubing without the need for additional connectors. The proposed microfluidic layout is illustrated in Figure 8. A secondary version with two channels was also created for future iterations.

The sealing is achieved mechanically using a PMMA holder, which allows the assembly to be easily disassembled for cleaning and reuse (see Figure 9). Fluidic connections are made using a 1/16" tubing inserted directly into the chip's orifices.

In the following table the specs of the microfluidic system are summarized.

Table 2: Specs of the microfluidics.

Microfluidic System	
Material	PDMS
Flow requirements	10-50 µL/min
Number of channels	8 (for Design 1) 2 (for Design 2)
Width of the channels	250 µm
Height of the channels	100 µm
Footprint	341 mm ²
Operating temperature range	NA
Mechanical or permanent sealing	Mechanical



The modification in the microfluidic chip required to make minor changes in the cover of the holder (see Figure 11), adapting the inlet and the outlet in the PMMA lid.




Figure 11. Adapted holder for the microfluidic chip design containing the single vertical channel.

3.3. InP active elements and interface with the SiN platform

3.3.1. InP active elements

Coupons design and specifications

SMART has developed two separate fabrication flows for 1) Semiconductor optical amplifiers (SOA's) and 2) uni-travelling carrier photodetectors (UTC-PDs). In the first flow, for SOA's, SMART uses a waveguide core with multi quantum wells (MQWs) as the active material. In the second flow, for UTC-PD's, SMART is using an InGaAsP core with a bandgap of 1.25 μm (Q1.25), plus an InGaAs layer right on top of it. In both flows, the previously described layer stacks are co-integrated with a passive waveguide core Q1.25 stack (WG). The detailed layer stacks are depicted in Figure 12.

After III-V fabrication, the under-etch of the sacrificial layers (InAlAs/InGaAs in Figure 12) is performed. The coupons are then released and supported only by photoresist tethers. These supporting tethers break when the coupons are transfer-printed by XCEL. This is schematically depicted in Figure 13. The transfer printing process is in progress.




Figure 12: (a) SOA layer stack (b) UTC layer stack (c) Passive WG layer stack.

The micro-transfer-printable-ready coupons mentioned above are derived from SMART's monolithic InP platform and are still under development. Consequently, the data provided below regarding the specifications of the active elements is only based on SMART's monolithic platform. This should be used as a reference, but it should be highlighted that critical changes such as a thinner n-InP cladding (to enable evanescent coupling), or a different substrate, might may an impact on device performance. The specifications are listed in Table 3, Table 4, and Table 5.




Figure 13: III-V integration and RDL process flow: a) spray coat BCB adhesive, b) μ TP, c) Remove resist tethers, d) spin coat BCB, e) open contact pads, f) sputter Cu seed layer, g) lithography, h) electroplating, i) etch back.

Table 3: SMART C-band uni-travelling carrier photodiode (UTC-PD) specifications in the monolithic platform.

BB	Parameters	Target specs	1 st run specs
C-band PD	Photodiode type	UTC	UTC
	Dark current (nA)	10 (typ.) - 20 μ m (max)	t.b.d.
	Responsivity (A/W)	0.8	t.b.d.
	Sensitivity (dBm)	<-50	t.b.d.
	Bandwidth (GHz)	30 (typ.) – 50 (max)	t.b.d.
	Coupling losses	1.5 dB	t.b.d.
	Footprint	2300x45 μ m ²	2300x45 μ m ²
	Mode	TM?	TE

Table 4: SMART C-band semiconductor optical amplifier (SOA) specifications for the monolithic platform.

BB	Parameters	Target specs	1 st run specs
C-band SOA	Wavelength band (nm)	1550	1550
	Linewidth (kHz)	$\cong 100$	t.b.d.
	Footprint (μm^2)	2500x45	2500x45
	Gain (dB)	10	13 @ 2kA/cm ² 22 @ 4kA/cm ²
	Temperature operating range	RT	RT
	Power consumption	100 mW	t.b.d.
	Fwd current (mA)	120 (typ.) – 200 (max)	t.b.d.
	Current stability (mA)	0.1	t.b.d.
	Power dissipation (mW)	10 (typ.) – 15 (max)	t.b.d.
	Mode	TE	TE

Table 5: SMART C-band low-linewidth tunable laser (LoLiTuLa) specifications for the monolithic platform.

BB	Parameters	Target specs	1 st run specs
C-band Tunable laser	Laser type	Tunable laser	Tunable laser
	Wavelength band/range (nm)	1550 ± 40	1550 ± 40
	Linewidth	$\cong 100$ kHz	t.b.d.
	Footprint (μm^2)	2500x400	2500x400
	Output power	10 dBm	10 dBm
	Temperature operating range	RT	RT
	Power consumption	100 mW	t.b.d.
	Fwd current (mA)	120 (typ.) – 200 (max)	t.b.d.
	Current stability (mA)	0.1	t.b.d.
	Power dissipation (mW)	10 (typ.) – 15 (max)	t.b.d.
Linear EOPM	Reverse bias (V)	0 (typ.) – 8 (max)	t.b.d.
	Bandwidth (nm)	30	t.b.d.
	Power dissipation (mW)	10 (typ.) – 50 (max)	t.b.d.
Left ring (EOPM)	Reverse bias	0 (typ.) – 8 (max)	t.b.d.
	Bandwidth	30	t.b.d.
	Power dissipation	10 (typ.) – 50 (max)	t.b.d.
Right ring (EOPM)	Reverse bias	0 (typ.) – 8 (max)	t.b.d.
	Bandwidth	30	t.b.d.
	Power dissipation	10 (typ.) – 50 (max)	t.b.d.

3.3.2. InP-to-SiN interface

The UOI undertook the design and optimization of the optical coupling interface between III-V active components and the SiN photonic platform. An adiabatic coupling scheme was developed to enable efficient light transfer between InP-based active devices—such as Semiconductor Optical Amplifiers (SOAs) and Uni-Traveling Carrier Photodiodes (UTC-PDs)—and the SiN waveguide platform, operating at 1550 nm for TE polarization. To address the significant refractive index mismatch between the SiN waveguide and the III-V epitaxial stack, a polycrystalline silicon (p-Si) intermediate layer with 350 nm thickness was introduced. A two-stage taper configuration was optimized: the first stage guides light from the SiN waveguide into the p-Si waveguide, while the second stage transitions the optical mode into one that overlaps effectively with the active region of the III-V stack—either for coupling into a UTC-PD or from an SOA-based laser, as illustrated in Figure 14a. This design enables efficient bidirectional coupling between the SiN platform and the III-V components integrated by micro-transfer printing. The final optimized designs achieved a **coupling efficiency of 94.8% for SOA-based laser structures**, with back-reflection maintained below the -50 dB threshold, and **58.2% optical absorption for UTC-PD structures** (see table in Figure 14b). The optimized parameters can be found in Figure 14a. The interface designs were fully aligned with the fabrication capabilities and constraints of the contributing partners: LGT (SiN and p-Si photonic platform), SMART (InP active devices), and XCEL (micro-transfer printing). Key input parameters included refractive indices, cross-sectional layer structures, individual layer thicknesses, critical fabrication dimensions, and alignment tolerances. The optimization procedure, along with the final design specifications and photonic mask layout for the first fabrication run, is comprehensively documented in Deliverable D3.1.

Characterization of the SiN-to-p-Si interface at AUTH verified that the optimal taper length is 100 μm . For this taper length and widths of 3 μm and 5 μm for the SiN and p-Si, respectively, the highest **experimental transmission** achieved was **94.1 %**. The **propagation losses** were consistent across the different variants of waveguide widths, with values ranging from **-8.79 dB/cm to -9.98 dB/cm**.




Figure 14: Design and optimization of InP-to-SiN interface. (a) Schematic illustration of the adiabatic coupling scheme between III-V active components and the SiN photonic platform, applicable to both SOA-based lasers and UTC-PDs. (b) Optimized total coupling efficiency, incorporating all losses due to waveguide widening and taper transitions.

Polarization rotator

The SOA and the tunable laser provided by SMART operate in TE polarization mode, whereas the plasmonic sensor requires TM polarization. To bridge this polarization mismatch, the UOI designed a polarization rotator that converts the light polarization from TE to TM. The design is based on a polarization rotation scheme proposed in earlier theoretical work², which employed an amorphous silicon (a-Si) overlayer on a silicon-on-insulator (SOI) platform. This concept was adapted to the LGT SiN platform by using a 350nm-thick p-Si overlayer—previously identified as optimal for the III-V to SiN interface. Utilizing the same p-Si layer allows both the polarization rotator and the III-V coupling interface to be defined in a single lithographic step, thus streamlining the fabrication process. Platform-specific constraints were carefully accounted for, including the thickness of the SiN and SiO₂ layers, as well as the taper tip geometry (Figure 15a). The optimized coupling configuration positions the taper tips outside the SiN waveguide core, as shown in the top view of Figure 15a. This "tip-out" geometry yields improved performance due to reduced scattering losses. Additionally, the optimal SiN and p-Si waveguide widths and polarization rotation length were defined to achieve a **high conversion efficiency, reaching 88%** (see table in Figure 15b). Tolerance and misalignment analyses were performed to assess the robustness and manufacturability of the design. It was found that fabrication imperfections and variations in refractive indices can be compensated through thermal tuning, requiring a temperature increase of only $\Delta T = 15$ K. Furthermore, back reflections in both TE and TM modes were kept below the -50 dB threshold. The optimization procedure, along with the final design specifications, is comprehensively presented in Deliverable D3.2.




Figure 15: Design and optimization of TE-to-TM polarization rotation. (a) Schematic illustration of the coupling scheme, including a cross-sectional view of the LGT photonic waveguide platform, and a top view of the optimized configuration. (b) Optimized design parameters of the proposed polarization rotator, intended for implementation in the second fabrication run of the AMBROSIA project.

² Y. Xiong, et al, "Robust Silicon Waveguide Polarization Rotator With an Amorphous Silicon Overlayer," in IEEE Photonics Journal, vol. 6, no. 2, pp. 1-8, April 2014, Art no. 2200308, doi: 10.1109/JPHOT.2014.2306827.

3.4. Photonic neuromorphic circuit

In the initial SiN fabrication run, key passive components essential for the MZI weighting elements were successfully fabricated for C-band operation, their performance is summarized in Table 6.

Table 6: Specifications and performance of SiN passive elements.

Passive LPCVD SiN platform (2μm SiO ₂ /400nm SiN/3.2μm SiO ₂ /Si substrate)			
Components	Parameters	TE - C band (1550nm)	TM - C band (1550nm)
Grating couplers	Coupling angle (degree)	15	25
	Central wavelength (nm)	1567	1581
	Max. Coupling efficiency (dB)	5.1	13
	1 dB bandwidth (nm)	59	64
Strip waveguide	Width (μm)	1.0	1.0
	Insertion loss (dB/cm)	0.35	0.35
Waveguide bend	Radius	60	60
	Insertion loss (dB/unit)	0.02	0.02
Waveguide cross	Insertion loss (dB/unit)	0.23	0.21
	Crosstalk (dB/unit)	-42	-44
MMI (1×2)	Insertion loss (dB/unit)	0.10	0.06
	Imbalance	0.62	n/a
MMI (2×2)	Insertion loss (dB/unit)	0.14	0.14
	Imbalance	0.44	n/a

As outlined in D3.2, two configurations for the weighting elements in the photonic neuromorphic circuit have been explored: one utilizing a UV-trimmed Mach-Zehnder Interferometer (MZI) combined with a tungsten filament, and the other employing a non-volatile, electro-refractive Phase Change Material (PCM)-based MZI. Both configurations were integrated within the silicon nitride (SiN) photonic Crossbar architecture illustrated in Figure 16.




Figure 16: Schematic of the neuromorphic crossbar architecture.

In the first-generation neuromorphic crossbar, the synaptic weights were implemented using a rapid, non-invasive UV-trimming process that enables adjusting the refractive index of the symmetric MZI arms to achieve the desired transmission weights. In this case, the UV weights were configured to form a 2-bit content-addressable memory (CAM) table with 9 different stored words. Each input search bit (SeB) requires 2 optically modulated signals, SeB and its inverse $\overline{\text{SeB}}$. The full 2-bit search word therefore requires 4 optically modulated signals, SeB1, $\overline{\text{SeB1}}$, SeB2 and $\overline{\text{SeB2}}$, as illustrated in Figure 17.




Figure 17: (a) graphic design system (GDS) of the crossbar waveguide layer and the corresponding weights of each MZI (b) CAM table of the UV weights in both the optical domain and the binary logic domain.

The UV weights were characterized by injecting into each of the 4 inputs and measuring the 9 outputs. Extinction ratios between the ON and OFF states ranged from 5-6 dB up to 20-25 dB. Figure 18 shows the normalized extinction ratios of each UV weight in the OFF state, where values near 0 dB indicate that the weight is still in the ON state. Ratios reaching up to 21 dB confirm the effectiveness of the UV writing process.




Figure 18: Extinction ratios of every weight that has been UV exposed.

In the second-generation neuromorphic crossbar, the weighting elements will be realised by exploiting the modulation effect of PCMs placed on top of the MZI arms, to enable a precise and reconfigurable control of the MZI response. The performance of the proposed approach has been evaluated on individual PCM switches that were included in the initial fabrication run. The weighting functionality of the MZI phase shifters was assessed by injecting a tunable continuous-wave laser (ranging

from 1520 nm to 1620 nm) into one of its inputs, applying a pulsed electrical signal to the ITO heater located on one of its arms, and then measuring their optical output using a photodetector. The applied electrical signal included both RESET pulses, with an amplitude of 13 V and a width of 60 μ s, to switch the material to its amorphous state, and SET pulses, with a lower voltage of 8 V and a width ranging from 3 to 9 μ s, to transition the material to intermediate levels of its crystalline state. These experiments enabled the estimation of the phase-shifting efficiency and optical losses of the PCM phase shifters in both amorphous and crystalline states, as summarized in Table 7.

Table 7: PCM phase-shifters insertion loss and efficiency.

Polarisation	MZI Loss (dB/unit)	Amorphous state (dB/ μ m)	Crystalline state (dB/ μ m)	Shifting Efficiency (π/μ m)
TE	1.5	0.008	0.021	0.008
TM	N/A	0.011	0.013	0.007

Initial cycling experiments, conducted on devices from the same fabrication batch, have enabled a comparative analysis of insertion loss between the amorphous and two distinct crystalline states of the PCM switches over 7 million switching cycles, as shown in Figure 19(a). This represents the current state of the art for these devices. Additional experiments have also enabled the generation of up to 32 distinct intermediate states between the crystalline and amorphous phases. Amongst them, 10 levels were identified as fully reliable, consistently producing stable results across 10 repeated switching cycles as shown in Figure 19(b).




Figure 19: (a) Insertion loss for the amorphous and two distinct crystalline states over 7,000,000 switching cycles and (b) Reliable usable levels across 10 repeated switching cycles.

3.5. Electronics

During the first 30 months of the project, it came out that, apart from the final printed circuit board (PCB), 3 separate PCBs are required to reach the project's goals: one for the MZIs' heaters drive, one for the neuromorphic chip drive and one for the control of the active elements (lasers, photodetectors). The technical specifications for the design and manufacture of the above PCBs are given below.

PCB for the MZI sensors' heater control

A dedicated electrical interface subsystem was specified and designed in order to control the phase shifters of MZI plasmo-phonic sensors. The central component of this subsystem is a custom PCB, fabricated from FR-4 material and architected to function as a passive probe card. The primary design specification for the chip-side interface mandates an array of ultra-fine tungsten pins, selected for their high durability and low contact resistance. This array is configured to precisely match the layout of the sensor's heater pads, which are specified as a 28-channel array with a pad size of $150\ \mu\text{m} \times 150\ \mu\text{m}$ and a pitch of $300\ \mu\text{m}$. For external connectivity, the system specifies a standard 40-pin edge connector to ensure a robust, high-channel-count interface with laboratory-grade current sources. The mechanical design specifies the integration of the probe card with a manually controlled X-Y-Z stage to achieve the high-precision alignment required for repeatable probing. A more detailed description, including mechanical drawings and operational procedures, can be found in Deliverable 4.2.

PCB for testing the neuromorphic chip

A custom printed circuit board (PCB) system was designed to provide a compact and easily configurable control system for automated testing of the Xbar chip, as well. Figure 20 illustrates the overall PCB architecture and the detailed pulse generator design.




Figure 20: (a) PCB architecture and IO and (b) Pulse generator/driver.

The overall system consists of two stacked PCBs, as shown in Figure 21. The top (right) board integrates all input/output interfaces, processing units, digital-to-analog converters (DACs), and DC amplification circuits. The bottom (left) board contains the high-speed pulse drivers and current measurement circuitry.




Figure 21: Overall PCB system.

Through the board interface, a computer can control both the PCM and thermo-optic modulators on the chip, while simultaneously interfacing with optical components, such as lasers and detectors. This capability enables the implementation of closed-loop control systems and automated feedback mechanisms, paving the way for fully automated chip characterization and testing.




Figure 22: PCB generated and measured (a) amorphization pulse, (b) crystallization pulse and (c) their optical response.

Initial testing of the PCB demonstrated that the system can generate pulses with suitable characteristics for the modulation of the PCM weights (Figure 22 (c)). Specifically, the PCB can produce high-energy pulses of up to 50 V with durations from DC to 100 ns. This is sufficient to drive the PCM into its amorphous state with 60 μ s pulses (Figure 22(a)), as well as promoting gradual recrystallization with 6-8 μ s pulses (Figure 22(b)). While full integration is still pending, these results highlight the potential of the system for precise and scalable control of PCM-based photonic devices.

PCB for driving the active elements

A separate PCB will be fabricated to speed up the characterization process of the active elements' chip of the 1st run. A complete electrical design, validated through simulation, has been established for this control board that will provide the stable, low-noise voltage and current sources required by the InP tunable lasers, SOAs, and photodetectors.

The hosting PCB is shown in Figure 23 and will include:

- 6 separate circuits for tunable laser control (remaining 2 lasers will have only sockets for connections to external devices)
- 6 separate circuits for SOA control (remaining 4 SOAs can be connected to external devices)
- 4 separate circuits for UTC PDs control (remaining 2 PDs will be possible sensing externally)
- SPI interface (for communication with ARDUINO controller)
- Multiple voltage sources for DAC operations




Figure 23: Schematic view of the PCB used for driving and controlling of each active component

Final PCB and assembly process

The assembly of the final demonstrator system is defined by a structured integration process flow. The process begins with the fabrication of the unified FR-4 PCB, followed by the population of all surface-mount electronic components, including DACs, operational amplifiers, drivers, and the Arduino microcontroller interface. The subsequent critical step is the mechanical and electrical integration of the photonic chips. The PICs will be affixed to the main PCB, and electrical connections will be established via gold wirebonding from the PIC's contact pads to the corresponding ENIG pads on the PCB. This process ensures short, direct electrical paths for optimal signal integrity. Once the PICs are integrated, the Arduino-compatible microcontroller board will be connected. The fully assembled electronic control module will then undergo a comprehensive series of power-up and functional verification tests before being integrated into the final demonstrator chassis. This step-wise process flow is designed to ensure reliability and facilitate systematic testing at each stage of integration.

3.6. Read-out system

The AMBROSIA read-out system is the complete electronic and software framework designed to control the integrated photonic components and translate their optical signals into a final diagnostic result. The system's architecture is built around a central microcontroller that intelligently manages a suite of custom-designed Printed Circuit Boards (PCBs). Each PCB provides the specialized electrical driving,

biasing, and signal conditioning required by a specific part of the photonic platform, from the neuromorphic processor to the active InP components.

The hardware foundation of the read-out system currently consists of three distinct, validated subsystems. First, a high-voltage, high-speed pulse and DC driver board provides comprehensive control of the neuromorphic chip, capable of programming both the PCM-based and thermo-optic modulators. Second, a dedicated probe card subsystem enables precise thermal tuning of the plasmo-photonic sensors via their on-chip heaters. Finally, a complete electrical design, validated through simulation, has been established for a control board that provides the stable, low-noise voltage and current sources required by the InP tunable lasers, SOAs, and photodetectors.

At the core of the control logic is an Arduino-compatible microcontroller, which serves as the master controller for the entire system. Communication between the Arduino and the various peripherals on the custom PCBs, such as the Digital-to-Analog Converters (DACs) and Analog-to-Digital Converters (ADCs), is implemented using a Serial Peripheral Interface (SPI) bus. This standardized bus architecture simplifies the hardware layout and ensures a robust, high-speed data link for sending commands and receiving data, allowing for deterministic control of each analog channel.

The operational control of the read-out system will be managed by custom firmware and application software. The firmware, running on the Arduino, will be responsible for executing low-level commands, such as writing specific values to DACs via SPI transactions. The high-level application software will provide the user interface (UI) for managing the diagnostic workflow. A critical component of this software will be a dedicated diagnostics and testing page, enabling manual control over individual hardware components for system verification and characterization.

The final phase of development for the read-out system, involves the consolidation of these proven, individual circuit designs onto a single, unified PCB. This integrated board, once fabricated and assembled, will represent the complete electronic heart of the AMBROSIA demonstrator. In conjunction with the corresponding software adaptation, this final hardware will form the fully assembled electronic read-out and diagnostic system, ready for comprehensive system-level testing and validation.

3.7. PoC platform

The AMBROSIA project aims to develop a portable, integrated Point-of-Care (PoC) diagnostic platform for the rapid and intelligent diagnosis of sepsis. The system is based on plasmo-photonic biosensors and is designed to be compact, affordable, and suitable for clinical environments. The first version of the PoC prototype, developed by BIOPIX-T in collaboration with several partners, integrates multiple subsystems into a cohesive diagnostic unit. A detailed presentation of the working prototype unit (version 1.0) is described in Deliverable 5.1, while a brief overview of the constituent components is given below:

1. Sensor & Microfluidics:

The core of the system includes a sensor and microfluidic subsystem that uses a PDMS-based chip with eight channels to guide biological samples over the sensor surface. Two types of sensors are considered: a balanced Mach-Zehnder Interferometer and a bimodal plasmo-photonic interferometer, both designed for high sensitivity and specificity. The microfluidic design ensures precise sample delivery and waste removal, with mechanical sealing and compatibility with standard tubing.

2. Laser/Photodiodes:

In the final prototype, the sensor chip will be equipped with integrated InP lasers and photodiodes developed by SMART and printed on the chip via μ TP by XCEL. In version 1.0, external active components have been used to allow for calibration flexibility.

3. Microcontroller:

The microcontroller, based on an Arduino platform, orchestrates the operation of all subsystems, including the micropumps, temperature control, and data acquisition.

4. Micropumps:

Fluid movement is managed by Bartels mp6-liq piezoelectric micropumps, which are electronically controlled and calibrated for flow rates between 5 and 40 microliters per minute. These pumps are supported by passive check valves and a dedicated driver module, ensuring precise and reproducible fluid handling. A calibration curve was established to fine-tune flow rates based on voltage input.

5. Thermal Management:

Thermal regulation is achieved using a Peltier element and a DS18B20 temperature sensor, which monitors internal conditions and displays real-time data on an OLED screen. Cooling is supported by two small fans controlled via transistors, which are activated automatically when the temperature exceeds 27°C.

6. Power Supply:

The entire system is powered through a standard DC input (5.5mm x 2.5mm jack).

7. Software & GUI

The software interface includes a desktop application that connects to the device via a USB port. It enables users to initiate tests, monitor system status, and visualize diagnostic results in real-time. Data is stored in JSON format and can be exported as CSV for further analysis. The application is designed to be user-friendly and supports the validation of analytical protocols.

8. Mechanical Enclosure

Mechanically, the prototype is housed in a 3D-printed enclosure made from PLA, with dimensions of 30 cm $\times$ 15 cm $\times$ 20 cm. The design includes separate compartments for electronics, sample input, sensor placement, and waste collection, ensuring ease of use and maintenance, as shown in D2.1.

Figure 24 shows a photograph of the connected components of version 1.0 during a validation experiment at AUTH premises. The operational scheme in Figure 24 follows a precise sequence: the sample is introduced via vials, guided through tubes to the sensor chip, analyzed in real-time, and disposed of through the pump into the waste compartment. The results are visualized in a desktop application, as shown in Figure 25.

Looking ahead, version 2.0 of the PoC platform will incorporate finalized components from project partners, including an upgraded printed circuit board with a neuromorphic chip for data classification. This version will feature improved integration, enhanced data analytics, and compliance with regulatory standards, paving the way for clinical validation and eventual commercialization.




Figure 24: Experimental setup of the subsystems of the PoC device.



Figure 25: Desktop application showing the results of the experiment.

4. Use cases and Artificial intelligence training models for sepsis detection and classification

4.1. Sepsis scenarios and targeted set of biomarkers and pathogens

The definition of use cases, target biomarkers, pathogens, and assay parameters in AMBROSIA has been established in the first stage of the project, based on clinically validated data and well-documented diagnostic criteria. The selected biomarkers and pathogens reflect current clinical standards for the early detection and management of sepsis, and their relevance has been confirmed through a thorough review included in Deliverables D1.1, D2.1 and D1.2. All associated assay ranges, clinical scenarios, and system requirements—including sample type, volume specifications for different patient groups, and a maximum

analysis time of 30 minutes— were defined accordingly and remain unchanged. The following tables summarize the sepsis scenarios and relevant clinical data.

Table 8: List of sepsis biomarkers and bacteria along with correlation to the SOFA score.

	Inflammation	Sepsis	Septic Shock
SOFA	---	≥2	≥4
CRP	< 50 mg/dL	≥50 mg/dL	≥50 mg/dL
IL6	< 1 pg/mL	≥ 1 pg/mL	≥ 1 pg/mL
PCT	< 0.5 ng/mL	>0.5 ng/mL	> 0.5 ng/mL
MR-proADM	>1.5 nM	>1.5 nM	> 2 nM
<i>E. coli</i>	-	10-10 ⁴ cfu/mL	10-10 ⁴ cfu/mL
<i>S. aureus</i>	-	10-10 ⁴ cfu/mL	10-10 ⁴ cfu/mL
<i>P. aeruginosa</i>	-	10-10 ⁴ cfu/mL	10-10 ⁴ cfu/mL

Table 9: List of sepsis biomarkers and bacteria selected and the assay range requirements for different clinical scenarios

Biomarker	Concentrations	
CRP	Assay range	<0.2 to >20 mg/dL
	Normal level healthy subjects	< 5 mg/dL
	Non-infectious SIRS	10 to 40 mg/dL
	Bacterial infection	40 to >200 mg/dL
	Sepsis/septic shock	170 to > 200 mg/dL
IL-6	Assay range	5 to 1000 pg/mL
	Healthy subjects	< 5 pg/mL
	Non-infectious SIRS	40-80 pg/mL
	Sepsis/septic shock	> 120 pg/mL
PCT	Assay range	0.02 to 75 ng/mL
	Healthy subjects	<0.1 ng/mL
	Bacterial infection	<0.1 to > 75 ng/mL
	Sepsis/septic shock	> 0.5 ng/mL
MR-pro-ADM	Assay range	0.1 to 20 nM
	Healthy subjects	< 0.5 nM
	Non-infectious SIRS	0.6- 1.1. nM
	Sepsis/septic shock	> 1.1 nM/L
Bacteria	Specie	Concentration
	<i>Escherichia coli</i>	10-10 ⁴ cfu/mL
	<i>Pseudomonas aeruginosa</i>	10-10 ⁴ cfu/mL
	<i>Staphylococcus aureus</i>	10-10 ⁴ cfu/mL

4.2. Sepsis diagnosis and classification data

In 2023 and 2024, 1.730 patients were selected for the AI analysis. All patients were activated by the Vall d'Hebron Hospital Sepsis Code and were donors to the Sepsis Bank of the Vall d'Hebron University Hospital Biobank. The data provided included gender, age, sepsis diagnosis type 2, sepsis diagnosis type 3, positive vs. negative blood culture, positive microorganism in the blood culture, lactate (nmol/L), C-reactive protein (mg/dL), procalcitonin (ng/mL), interleukin 6 (pg/mL), MR proADM (nmol/mL), global SOFA score, cardiovascular SOFA score, respiratory SOFA score, hepatic SOFA score, renal SOFA score, and neurologic SOFA score.

The patients were classified according to the Sepsis 2 and 3 definitions. According to the sepsis 2 definitions, the categories of septic shock (705 patients), severe sepsis (662 patients), infection without Systemic Inflammatory Response Syndrome (SIRS) (90 patients), no infection-SIRS (227 patients) and no infection-no SIRS (46 patients) were established. Following the sepsis 3 classifications, the categories of septic shock (438 patients), sepsis (906 patients), infection (33 patients), no infection (271 patients) and complicated infection (67 patients) were established. Fifteen patients could not be classified according to the sepsis 3 definitions because lactates were not determined. Patients with complicated infection correspond to patients who have presented an infection that has progressed with clinical severity but do not meet the criteria of the sepsis 3 definitions to be classified as sepsis because the SOFA has been less than 2 points or who cannot be classified as septic shock because although they have required vasoactive support they have not presented hyperlactatemia.

The above real-world datasets aim to help build reliable machine learning models for the early detection and stratification of sepsis and were submitted to AUTH in the form of an Excel database. A structured description of the datasets used, categorized by their development stage and clinical enhancements is given below. Because the datasets used in Deep Learning models are crucial for their predictive ability, the curation of each iteration of the dataset is achieved with the aid of a proposed methodology named Domain-Guided Data Refinement, a data-driven technique to enable better neural networks.

Version 1: Initial Dataset (D_1)

The first version of the dataset (D_1) was developed as a foundational resource to train machine learning models for detecting sepsis and differentiating it from non-septic conditions. It includes:

- 785 patients, with collected variables such as:
 - Basic demographic data
 - SOFA scores
 - Blood culture results
 - Microorganism identification (when available)

These clinical indicators (e.g., SOFA, cultures) often take hours to return, making them unreliable for early detection scenarios.

Categories in D_1 :

- Non-Infectious SIRS: 175 samples
- Sepsis: 242 samples
- Septic Shock: 331 samples

Biomarkers Available:

- C-Reactive Protein (CRP): Indicates systemic inflammation.
- Interleukin-6 (IL-6): A pro-inflammatory cytokine elevated in sepsis.
- Procalcitonin (PCT): A rapid biomarker for bacterial infection and systemic inflammation.

Version 2: Refined Dataset with Expert Labels (D_2)

Version 2 (D_2) represents a significant update following a *Domain-Guided Data Review (DGDR)* process. Domain experts and physicians re-annotated the data and harmonized labels using three classification strategies (internal clinical labels, Sepsis-2, and Sepsis-3). This version includes:

- 949 patients, including all from D_1 and newly collected cases.
- Expanded biomarker panel to improve biological interpretability:
 - Lactate: Reflects tissue hypoxia and correlates with poor perfusion.
 - Mid-regional Pro-adrenomedullin (MR-proADM): Available for 585 patients; correlates with vascular integrity and sepsis severity.

Classification Types:

1. Hospital Internal Labels

- Infection without SIRS: 52
- No Infection: 31
- SIRS (Non-infectious): 130
- Severe Sepsis: 331
- Septic Shock: 404

2. Sepsis-2 Definitions

- Similar classes with slight terminological adjustments (e.g., “No infection - SIRS”)

3. Sepsis-3 Definitions

- Complicated Infection: 45
- Infection: 11
- No Infection: 159
- Sepsis: 489
- Septic Shock: 232

Version 3: Expanded Longitudinal Dataset with Yearly Breakdown (D_3)

Version 3 (D_3) includes ongoing and completed patient data across three years and reflects real-world hospital deployment. This version captures a broader spectrum of clinical progression and more representative patient diversity, that will total 2041 patients.

Substantial efforts were made to include more female patients in the dataset to lead to more fair predictive models.

Table 10: Version 3 Dataset.

Year	Total Patients	Female (%)	Infection without SIRS	No infection	Severe Sepsis	Septic Shock
2023 (Completed)	1030	375 (36.4%)	56 (5.4%)	188 (18.2%)	358 (34.7%)	428 (41.5%)
2024 (Ongoing)	756	316 (41.8%)	36 (4.8%)	116 (15.3%)	313 (41.4%)	291 (38.4%)
2025 (Ongoing)	255	109 (42.7%)	17 (6.6%)	49 (19.2%)	97 (38.0%)	92 (36.0%)

4.3. AI training models

Optics-Informed Neural Networks: converging hardware with training models

This section presents the integration of Deep Learning (DL) models with neuromorphic photonic hardware for ultra-fast and energy-efficient inference, focusing on sepsis detection. Photonic ANNs use experimentally derived activation functions—such as sigmoid, sinusoidal, and tanh—based on real hardware responses (e.g., WDM, MZM), and are modeled during training to reflect physical behavior. Additional activation functions, including ReSin, ExpSin, and Double Sine Squared, are implemented through programmable analog opto-electronic (OE) and opto-electro-optic (OEO) circuits. These are mathematically defined using variants of tanh and sine-squared functions, configured via transimpedance amplifiers and MZMs. Due to their narrow activation range, such functions are prone to vanishing gradients, requiring tailored training strategies and initialization methods that preserve variance and account for data distribution. Novel activation-aware and activation-agnostic initialization schemes can be employed to address these challenges and enable effective training aligned with hardware constraints.

Noise-aware Training in Optics-Informed Neural Networks

In order to enable noise-aware training in optics-informed neural networks, hardware-specific limitations—such as shot noise, thermal noise, weight read noise, limited bandwidth, and extinction ratio—are modeled as Additive White Gaussian Noise (AWGN) and incorporated into the training process. This approach leverages the intrinsic ability of neural networks to tolerate noise when properly trained to withstand it. Noise is introduced at different stages of the neuron’s operation, including the input, weight multiplication, and activation output, enhancing robustness during deployment. Additionally, an advanced initialization method is used, which iteratively estimates the optimal variance per layer by introducing a learnable scaling factor for each layer. This factor is optimized through an auxiliary classification task, while network weights remain fixed, and a regularization term, $\Omega(a)$ penalizes saturation outside the activation range. Once optimized, the scaling factors are used to re-initialize the network weights, and the model proceeds with standard backpropagation. This process enables the network to effectively compensate for hardware-induced noise and operational constraints, improving performance in real-world photonic implementations.

Training with Quantized and Mixed Precision Neural Networks

In order to enable training with quantized and mixed precision neural networks in photonic systems, it is essential to account for the analog precision degradation introduced by components such as ADCs, DACs, modulators, photodiodes, and amplifiers. These components contribute noise that reduces the effective bit resolution of the electro-optic link, especially at higher line rates. Simulating this degradation during training -typically through quantization- allows models to learn more robust representations, improving performance during deployment. Quantization transforms continuous values into discrete ones by rounding or truncating them within a defined floating-point range. While traditional quantization methods in deep learning focus on over-parameterized CNNs and often ignore input and bias quantization, photonic neural networks require tailored approaches due to their high-slope activations and physical limitations.

Photonic systems commonly use simple uniform quantization, differing from trainable and mixed-precision schemes typically used in software. To address this, post-training quantization methods calibrate pre-trained models to limited precision, while quantization-aware training incorporates bit constraints during the learning process, yielding significantly better performance. One such quantization-aware framework applies uniform quantization to all parameters during the forward pass, while ignoring

rounding during backpropagation. It adjusts the model to lower-precision computations by propagating the quantization error through the network. The quantization process involves scaling and shifting values based on the bit resolution and signal range, followed by dequantization to simulate analog behavior.

To minimize quantization error, the training dynamically updates proxy minimum and maximum values of signals based on their evolving distribution, using exponential averaging of their mean and variance over time. This normalized quantization process improves accuracy and can also be used for post-training calibration. Furthermore, advanced methods introduce stochastic mixed precision training, which gradually adjusts the bit resolution across layers based on their sensitivity to noise and their location within the architecture. This dynamic precision strategy significantly reduces inference time while maintaining high model accuracy, making it well-suited for deployment on photonic hardware.

Neural Network Architecture and Learning Objective

Let $\mathbf{x}_i \in \mathbb{R}^D$ be an input sample of dimensionality D , and let $\mathbf{y}_i \in \{0, 1\}^C$ be its one-hot encoded label vector, such that:

$$\sum_{c=1} y_{i,c} = 1$$

This dataset consists of N training examples spanning C distinct classes.

The model is a deep neural network with L **fully connected layers**, where the l -th layer contains N_l **neurons**, and the final (output) layer has C **neurons**, one per class. The network $f(\cdot) \in \mathbb{R}^C$ is trained to map input samples $\mathbf{x}_i$ to their corresponding one-hot labels $\mathbf{y}_i$, so that the output of the network provides predictions for in-domain samples. The activation function is denoted as σ , where $\sigma(\cdot)$ is typically a ReLU activation function for hidden layers and a softmax activation for the output layer. The output of the l -th layer can be calculated as:

$$\mathbf{z}_l = \mathbf{W}_l \sigma(\mathbf{z}_{l-1}) + \mathbf{b}_l, \quad \text{for } l = 1, \dots, L,$$

where $\mathbf{W}_l \in \mathbb{R}^{N_l \times N_{l-1}}$ is the weight matrix at the l -th layer, and $\mathbf{b}_l \in \mathbb{R}^{N_l}$ is the bias vector at the l -th layer. For the first layer, $\mathbf{z}_0 = \mathbf{x}_i$.

The final output layer $\hat{\mathbf{y}}_{i,c} \in \mathbb{R}^C$ represents the predicted class probabilities, calculated using the softmax function, which ensures that the sum of the predicted probabilities across all classes equals 1.

The network is trained by minimizing the cross-entropy loss between the predicted class probabilities $\hat{\mathbf{y}}_i$ and the true onehot label vector $\mathbf{y}_i$:

$$L_{class} = - \sum_{c=1} y_{i,c} \log(\hat{y}_{i,c})$$

where $y_{i,c}$ is the target label for class c and $\hat{y}_{i,c}$ is the predicted probability for class c . The network is also regularized with L_2 regularization on the weights in order to avoid overfitting, with a hyper-parameter λ controlling the strength of the regularization. Thus, the final learning objective L used to train the network is the sum of the classification loss and the regularization loss:

$$L = L_{class} + L_{reg}$$

The weights $\mathbf{W}_1, \mathbf{W}_2, \dots, \mathbf{W}_L$ and the biases $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_L$ are updated using gradient descent to minimize the total loss function, ensuring that the network learns to accurately map input samples to their corresponding labels.

Domain-Guided Data Refinement

To improve data quality and label reliability, we propose a **Domain-Guided Data Refinement (DGDR)** method. This iterative process improves the dataset $\mathcal{D}_i$ (starting from $\mathcal{D}_1$) through model prediction and expert validation.

We propose a domain-guided data refinement approach to improve the quality of a dataset $\mathcal{D}_i$ (initially $\mathcal{D}_i = \mathcal{D}_1$). The steps involved are as follows:

- 1) First, we train a deep neural network on the dataset $\mathcal{D}_i$ to predict labels for all samples. Let the input sample $\mathbf{x}_i \in \mathcal{D}_i$ have a corresponding predicted label $\hat{y}_i$ generated by the model. The prediction for each sample is computed as:

$$\hat{y}_i = f(\mathbf{x}_i), \forall \mathbf{x}_i \in \mathcal{D}_i$$

- 2) Next, we calculate feature importance scores for each feature of each sample in $\mathcal{D}_i$:

$$s_{i,j}^{\wedge} = \left| \frac{\partial L(\hat{y}_i)}{\partial x_{i,j}} \right|,$$

where $s_{i,j}$ denotes the importance score for the j -th feature of the i -th sample, $L(\hat{y}_i)$ is the loss function, and $x_{i,j}$ is the value of the j -th feature of the i -th sample. The normalized feature importance scores $s_{i,j}^{norm}$ are computed as:

$$s_{i,j}^{norm} = \frac{s_{i,j}}{\sum_{j=1}^D s_{i,j}}.$$

- 3) For each incorrectly classified sample $\mathbf{x}_i$, we focus on the top F most important features, based on the normalized feature importance scores. For these top F features, we check whether their values comply with predefined domain-specific guidelines. If any of the top F features of a sample do not meet the guidelines, the sample is flagged as inconsistent with domain knowledge (we discuss later how this check can be performed in a practical and efficient way). This can be expressed as: $\text{Flag}(\mathbf{x}_{i,j}, l_i) = 1$, if $f_b^D(\mathbf{x}_{i,j}, l_i) > 0$ and $\text{Flag}(\mathbf{x}_{i,j}, l_i) = 0$, otherwise, where $f_b^D(\mathbf{x}_{i,j}, l_i)$ encoded the prior domain knowledge and represents a function that checks whether a sample classified as l_i is biologically/physically plausible when biomarker b has the value of feature $x_{i,j}$.
- 4) After consulting with domain experts, the flagged samples are reviewed, and corrections are made based on domain knowledge. This can result to conclusions on which samples need correcting and whether or not new features or samples should be added to a new dataset $\mathcal{D}_{i+1}$. The data refinement process can be repeated iteratively by retraining the model on the updated dataset $\mathcal{D}_{i+1}$, recalculating feature importance scores, and refining the data further.

Note that domain knowledge is encoded in a set of functions $f_b^D(\mathbf{x}, l)$, where b denotes a biomarker, $\mathbf{x}$ the corresponding measurements and l the label predicted by a predictive model. Each $f_b^D(\mathbf{x}, l)$ is defined as follows: $f_b^D(\mathbf{x}, l) = 1$, if the measurement $\mathbf{x}$ of biomarker b cannot correspond to label l and: $f_b^D(\mathbf{x}, l) = 0$, otherwise. The aim of $f_b^D(\mathbf{x}, l)$ is not to reclassify the samples but rather to identify cases where the prediction provided by a predictive model is inconsistent with prior knowledge.

Experimental Evaluation

Identifying inconsistent samples with DGDR

Inconsistencies identified by the proposed approach are illustrated in Figure 26 using samples drawn from the test dataset. For instance, in sample ID “2”, the patient has CRP = 0.4 mg/dL, PCT = 0.22 ng/mL, and IL-6 = 12.51 pg/mL. In this case, it is impossible to determine whether the patient is at risk of developing Septic Shock. Therefore, the need for additional biomarkers arises. Furthermore, examples such as the patient with ID “8” exhibit high values across all biomarkers CRP = 44.1 mg/dL, PCT = 17.14 ng/mL, and IL-6 = 1894 pg/mL, yet it cannot be concluded that the patient has Non-infectious SIRS, highlighting the need for corrections in annotation.

Sample ID	CRP (mg/dL)	PCT (ng/mL)	IL-6 (pg/mL)	Label	Prediction
1	5.5	0.28	11.51	Septic Shock	Non-infectious SIRS
2	0.4	0.22	12.51	Septic Shock	Non-infectious SIRS
3	3.4	0.30	23.40	Septic Shock	Non-infectious SIRS
4	3.6	0.08	24.20	Sepsis	Non-infectious SIRS
5	7.4	0.09	27.20	Sepsis	Non-infectious SIRS
6	0.8	0.05	3.59	Sepsis	Non-infectious SIRS
7	2.6	0.04	27.76	Sepsis	Non-infectious SIRS
8	44.1	17.14	1894	Non-infectious SIRS	Septic Shock
9	12.7	9.01	139	Non-infectious SIRS	Septic Shock
10	13.3	46.69	2639	Non-infectious SIRS	Septic Shock

Figure 26: Examples of inconsistent samples identified by DGDR. Column ‘Label’ indicates the ground-truth label provided by physicians and prediction label indicates the output of the trained DNN.

Evaluation Metrics

Definitions of the evaluation metrics are given below:

- **Accuracy:** Of all the instances, how many did the model correctly predict?
- **Precision:** Of all the instances that the model predicted as positive, how many were actually positive?
- **Recall:** Of all the actual positive instances, how many did the model correctly predict as positive?
- **F1 Score:** What is the harmonic mean of precision and recall, balancing both metrics?

		POSITIVE	NEGATIVE
ACTUAL VALUES	POSITIVE	TP	FN
	NEGATIVE	FP	TN

$$Precision = \frac{TP}{TP + FP} \quad Recall = \frac{TP}{TP + FN}$$

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN}$$

$$F1\ Score = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

Results of DGDR

We evaluated the performance of a deep neural network (DNN) model on two different datasets for sepsis prediction tasks and compared it with a suite of classical machine learning models, including Decision

Trees, K-Nearest Neighbors (KNN), Logistic Regression, Support Vector Machines (SVM), and Ridge Classifiers. The experiments were performed over 10 independent runs and mean $\pm$ standard deviation metrics for accuracy and F-score are reported.

Superiority of the Photonic Neural Network

Across both datasets and all evaluation settings, the proposed DNN consistently outperformed all baseline models:

- On Dataset 1, which involved a binary classification task using CRP, PCT, and IL-6 biomarkers:
 - The DNN achieved an accuracy of 0.7711 ± 0.0268 and an F-score of 0.7049 ± 0.0550 .
 - These metrics were significantly higher than all other models. For instance, the best classical model (Logistic Regression, SVM, and Ridge Classifier) achieved an accuracy of 0.7604 ± 0.0261 and an F-score of only 0.6572 ± 0.0353 .
 - This demonstrates a relative improvement of approximately 1.4% in accuracy and 7.3% in F-score over the best baseline.
- On Dataset 2, which used an expanded biomarker set (Lactate, CRP, PCT, IL-6, and MR-proADM) and adopted the Sepsis-3 definition for a three-class classification:
 - The DNN reached an accuracy of 0.8006 ± 0.0274 and an F-score of 0.7412 ± 0.0503 .
 - Again, it outperformed all classical models, with the second-best being Logistic Regression at 0.7757 ± 0.0296 accuracy and 0.7056 ± 0.0368 F-score.
 - This yields a 3.2% gain in F-score and a 2.5% gain in accuracy over the closest competitor.

Impact of Dataset Composition and Labeling Scheme

Two key factors significantly influenced model performance:

1. Sepsis-III Label Definition and Addition of Biomarkers (Dataset 1 $\rightarrow$ Dataset 2):
 - In Figure 27, where the classification schema is kept constant, but different biomarker sets are used:
 - DNN accuracy increased from 0.7711 to 0.8006, and F-score from 0.7049 to 0.7412.
 - This reflects a 3.8% improvement in accuracy and 5.1% in F-score, suggesting that the Sepsis-III definition as well as the addition of biomarkers used yields improved predictive capabilities by the neural network.
2. Sepsis-III Label Definition and Addition of Biomarkers (Dataset 1 $\rightarrow$ Dataset 2):
 - In Figure 28, where the classification schema is kept constant, but different biomarker sets are used:

- For the DNN, performance improved from 0.5040 ± 0.0332 accuracy and 0.4900 ± 0.0312 F-score (Dataset 1) to 0.6245 ± 0.0299 accuracy and 0.5811 ± 0.0317 F-score (Dataset 2).
 - This marks a relative increase of 24% in accuracy and 18.6% in F-score, underscoring that the enriched biomarker set offers more informative features, boosting predictive performance.
3. Comparison of Hospital Label Definition and Sepsis-III Label Definition (Dataset 2):
- In Figure 29, where the biomarkers used are identical but the label definition changes from the hospital definition to the Sepsis-III label definition:
 - DNN accuracy increased from 0.5253 to 0.6090, and F-score from 0.4676 to 0.5690.
 - This reflects a 8.3% improvement in accuracy and 10.1% in F-score, suggesting that the Sepsis-III definition yields better-aligned outcome labels, enhancing the correlation between patient data and clinical outcomes.

Dataset 1 Experiments

TABLE Sepsis prediction evaluation on $\mathcal{D}_1$ using CRP, PCT and IL6 biomarkers. Each sample is classified to one of the following classes: “Non-Infectious SIRS” and “Sepsis or Septic Shock”. Mean and standard deviation of accuracy and F-score are shown, calculated over 10 runs.

Model	Accuracy	F-score
Decision Tree	0.6964 ± 0.0264	0.6989 ± 0.0320
KNN	0.7364 ± 0.0228	0.7055 ± 0.0318
Logistic Regression	0.7604 ± 0.0261	0.6572 ± 0.0353
SVM	0.7604 ± 0.0261	0.6572 ± 0.0353
Ridge Classifier	0.7604 ± 0.0261	0.6572 ± 0.0353
DNN	0.7711 ± 0.0268	0.7049 ± 0.0550

Dataset 2 Experiments

TABLE Sepsis (Sepsis-3 definition) prediction evaluation on $\mathcal{D}_2$ using Lactate, CRP, PCT, IL-6 and MR-proADM on predicting “Sepsis or Septic Shock”, “No Sepsis and no Septic Shock”. Mean and standard deviation of accuracy and F-score are shown, calculated over 10 runs.

Model	Accuracy	F-score
Decision Tree	0.6971 ± 0.0379	0.7020 ± 0.0402
KNN	0.7296 ± 0.0281	0.7054 ± 0.0303
Logistic Regression	0.7757 ± 0.0296	0.7055 ± 0.0468
SVM	0.7769 ± 0.0284	0.6825 ± 0.0401
Ridge Classifier	0.7752 ± 0.0283	0.6788 ± 0.0384
DNN	0.8006 ± 0.0274	0.7412 ± 0.0503

Figure 27: Sepsis-III Label Definition and Addition of Biomarkers (Dataset 1 → Dataset 2).

Dataset 1 Experiments

TABLE Sepsis prediction evaluation on $\mathcal{D}_1$ using CRP, PCT and IL6 biomarkers. Each sample is classified to one of the following classes: “Non-Infectious SIRS”, “Sepsis” and “Septic Shock”. Mean and standard deviation of accuracy and F-score are shown, calculated over 10 runs.

Model	Accuracy	F-score
Decision Tree	0.4231 ± 0.0164	0.4233 ± 0.0160
KNN	0.4333 ± 0.0344	0.4365 ± 0.0312
Logistic Regression	0.4898 ± 0.0327	0.4724 ± 0.0484
SVM	0.4769 ± 0.0391	0.4604 ± 0.0537
Ridge Classifier	0.4747 ± 0.0365	0.4324 ± 0.0630
DNN	0.5040 ± 0.0332	0.4900 ± 0.0312

Dataset 2 Experiments

TABLE Sepsis (Sepsis-3 definition) prediction evaluation on $\mathcal{D}_2$ using Lactate, CRP, PCT, and IL-6 biomarkers. Mean and standard deviation of accuracy and F-score are shown, calculated over 10 runs.

Model	Accuracy	F-score
Decision Tree	0.5234 ± 0.0251	0.5232 ± 0.0303
KNN	0.5582 ± 0.0278	0.5401 ± 0.0314
Logistic Regression	0.5795 ± 0.0288	0.5060 ± 0.0389
SVM	0.5813 ± 0.0264	0.4942 ± 0.0318
Ridge Classifier	0.5681 ± 0.0289	0.4686 ± 0.0349
DNN	0.6245 ± 0.0299	0.5811 ± 0.0317

Figure 28: Sepsis-III Label Definition, Different Biomarkers (Dataset 1 → Dataset 2).

Sepsis-II Experiments

TABLE IV: Sepsis (hospital definition) prediction evaluation on $\mathcal{D}_2$ using Lactate, CRP, PCT, IL-6 and MR-proADM biomarkers. Mean and standard deviation of accuracy and F-score are shown, calculated over 10 runs.

Model	Accuracy	F-score
Decision Tree	0.4036 ± 0.0267	0.4021 ± 0.0282
KNN	0.4127 ± 0.0171	0.3891 ± 0.0144
Logistic Regression	0.5036 ± 0.0304	0.4445 ± 0.0318
SVM	0.4952 ± 0.0280	0.4352 ± 0.0302
Ridge Classifier	0.5042 ± 0.0296	0.4371 ± 0.0337
DNN	0.5253 ± 0.0289	0.4676 ± 0.0367

Sepsis-III Experiments

TABLE VIII: Sepsis (Sepsis-3 definition) prediction evaluation on $\mathcal{D}_2$ using Lactate, CRP, PCT, IL-6 and MR-proADM biomarkers. Mean and standard deviation of accuracy and F-score are shown, calculated over 10 runs.

Model	Accuracy	F-score
Decision Tree	0.5325 ± 0.0474	0.5267 ± 0.0450
KNN	0.5500 ± 0.0379	0.5267 ± 0.0367
Logistic Regression	0.5729 ± 0.0243	0.5043 ± 0.0310
SVM	0.5880 ± 0.0279	0.5085 ± 0.0312
Ridge Classifier	0.5861 ± 0.0313	0.4967 ± 0.0359
DNN	0.6090 ± 0.0210	0.5690 ± 0.0187

Figure 29: Comparison of Hospital Label Definition and Sepsis-III Label Definition (Dataset 2).

Gender Fairness in AI Models

The results demonstrate that resampling female samples to achieve a balanced gender ratio improves fairness in sepsis prediction models. Specifically, for the **level of sepsis prediction task**, increasing the female-to-male ratio from the default 0.3:1 to 1:1 significantly improved the F-score for females (from 61.34% to 70.68%), reducing the performance gap with males (74.88%). This shows that balancing the dataset mitigates gender bias without negatively impacting male performance. The F-score, which accounts for both precision and recall, is particularly sensitive to imbalances, making it a strong indicator of fairness.

Notably, the **sepsis detection task** showed no significant gender disparity in the default setting, suggesting that bias is more pronounced in predicting sepsis severity. The 1:1 ratio emerged as the optimal balance, as further oversampling females (2:1) did not yield additional gains. These findings highlight that resampling is an effective technique for addressing gender bias in clinical AI models, particularly for complex tasks like severity prediction, and should be considered in fairness-aware model development.

TABLE Sepsis detection and level of sepsis prediction results using a fairer dataset version. The mean and standard deviation of accuracy and f-score are shown, averaged over 10 runs. * indicates the default $\mathcal{D}_1$ dataset female to male ratio.

Task	Female to Male Ratio	Accuracy	F-score
Sepsis Detection (on females)	0.3*	0.4978 ± 0.0315	0.4874 ± 0.0354
Sepsis Detection (on males)	0.3*	0.5087 ± 0.0257	0.4872 ± 0.0316
Level of Sepsis Prediction (on females)	0.3*	0.7229 ± 0.0506	0.6134 ± 0.0645
Level of Sepsis Prediction (on males)	0.3*	0.7845 ± 0.0214	0.6942 ± 0.0352
Level of Sepsis Prediction (on females)	1	0.7469 ± 0.0568	0.7068 ± 0.0764
Level of Sepsis Prediction (on males)	1	0.7892 ± 0.0317	0.7488 ± 0.0316
Level of Sepsis Prediction (on females)	2	0.7225 ± 0.0506	0.6843 ± 0.0757
Level of Sepsis Prediction (on males)	2	0.7854 ± 0.0249	0.7503 ± 0.0354

To conclude, in the final version of the dataset ($\mathcal{D}_3$) by using the resampling method with a female to male ratio equal to 3, the model achieved an F-score on males equal to **72.54%** and an F-score on females equal to **71.95%**, thus leading to more **fair** predictions.

5. Conclusions

In this deliverable, the designs and specifications of the subsystems associated with the AMBROSIA PoC platform for sepsis detection and classification have been reported. More specifically, the performance of the core components, including sensors and microfluidics, has been verified experimentally showing sensitivities up to 13312 nm/RIU. The InP active elements have been fabricated and their printing on the SiN sensor chip is in progress. The interface between the two platforms has been verified experimentally showing coupling efficiency up to 94%. The first version of the photonic neuromorphic chip has also been fabricated and characterized successfully. Regarding electronics, three separate PCBs have been designed to support the verification process of the subsystems. The design of the final PCB is frequently updated according to experimental results and the feedback received from the respective partners. A first version of the PoC platform and the read-out software has been developed and verified experimentally. Finally, the deliverable described the developed artificial intelligence (AI) training models to be used for sepsis detection and severity classification utilizing photonic neural networks.