Toward AI-based classification of ground-based auroral emission spectra

G. Cessateur[1], M. Le Lain[2], H. Lamy[1], S. Lefèvre[2,3], M.G. Johnsen[4], M. Barthelemy[5], G. Bertrand[6]

[1] Royal Belgian Institute for Space Aeronomy, Brussels, Belgium [2] IRISA, UMR CNRS 6074, Université Bretagne Sud, Vannes, France
[3] UiT – The Arctic University of Norway, Tromsø, Norway [4] Tromsø Geophysical Observatory, Faculty of Science and Technology, UiT – The Arctic University of Norway
[5] Univ. Grenoble Alpes, CNRS, IPAG, 38000 Grenoble, France [6] STAROS, Antibes, France

Why automate auroral spectroscopy?

- Auroral spectra contain direct diagnostics of the **type and characteristic energy of precipitating particles**.
- **ASIS at Skibotn** has recorded auroral spectra since **October 2023**: 405–687 nm, spectral resolution $\sim$1.2 nm, **one spectrum every 30 s, every night**.
- The archive is already too large for systematic manual inspection. The original practical goal was simple: **automatically flag characteristic emissions**, especially Hα and other diagnostic lines.

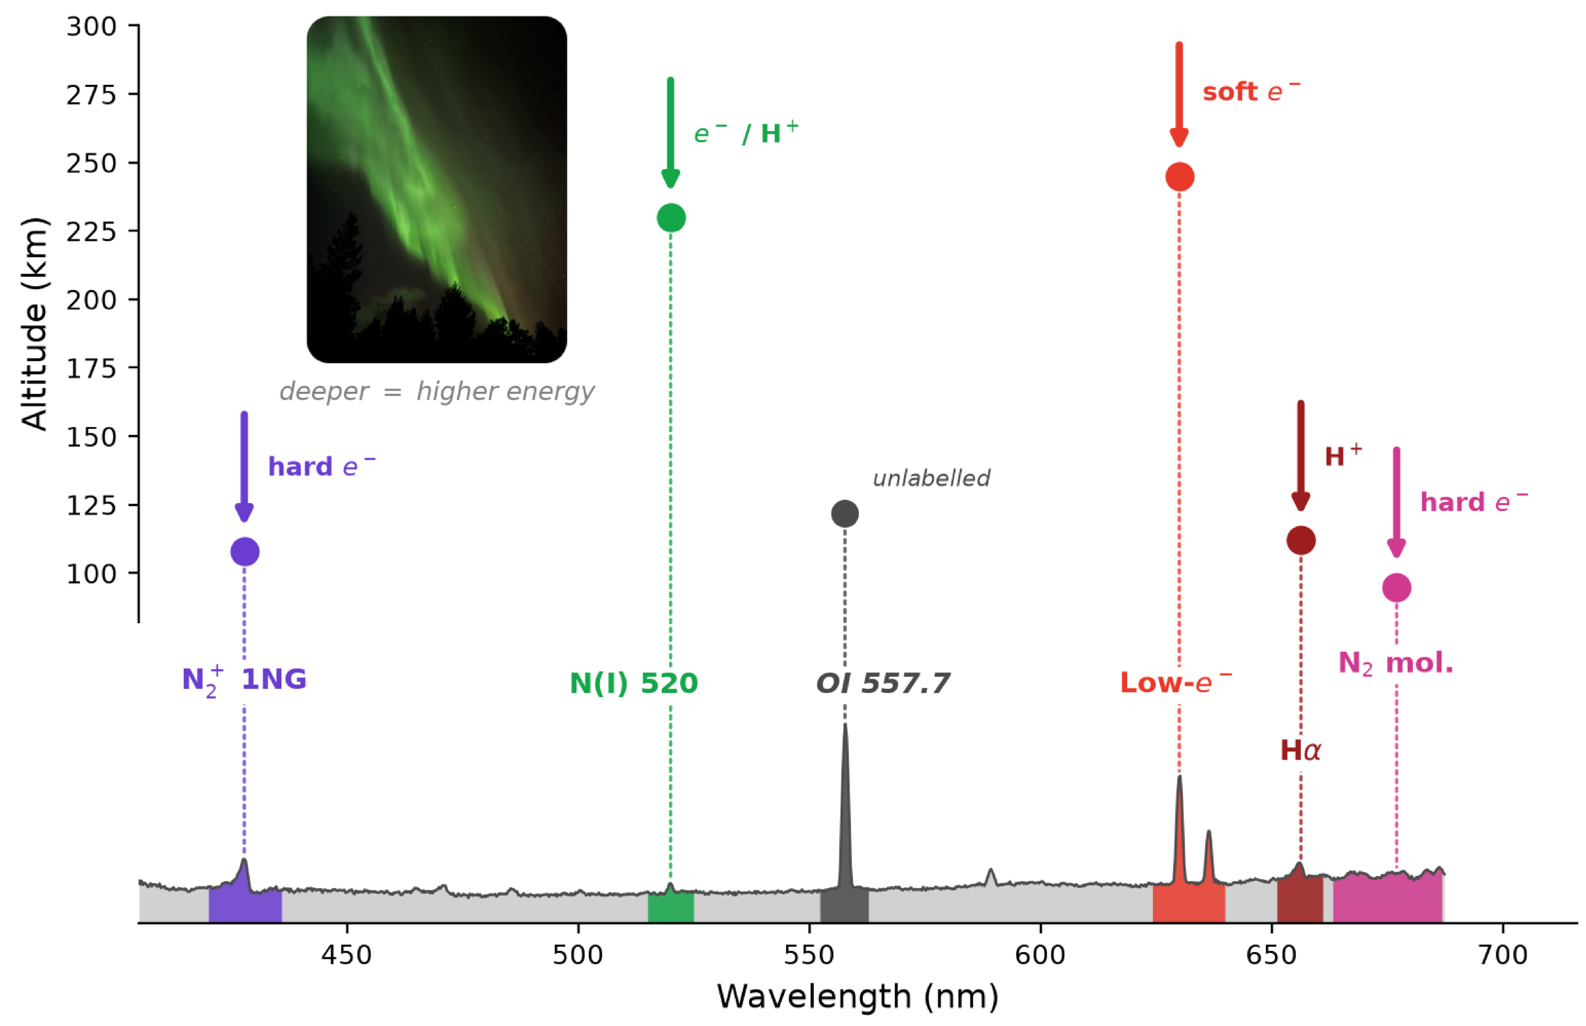

Figure: **From precipitation to spectral signature.** Different particle populations populate different emission altitudes and spectral features; proton precipitation produces Balmer emission such as Hα.

Can AI identify physically meaningful auroral spectral regimes at the scale of a multi-winter archive?

1. Supervised proof of concept: can AI reproduce expert labels?

- Published benchmark: **719 expert-annotated spectra**, with five possible physical labels per spectrum.
- Two supervised models were compared on the same spectra: a simple **Multi-Layer Perceptron (MLP)** baseline and a **1D Vision Transformer (ViT-1D)**.

1024-bin spectrum $\lambda = 405$–687 nm → **8 spectral patches** 128 bins each → **ViT encoder** self-attention across patches → **Classification head** 5 output labels

Label	Physical tracer	n
Hα	Proton precipitation	346
N_2 molecular	High-energy e^- precipitation	114
Low-e^-	Low-energy e^- precipitation	214
N(I) 520 nm	Atomic nitrogen emission	307
N_2^+ 1NG	N_2 ionisation / energetic e^-	267

Average Precision (AP), averaged over the five classes:
MLP: **81.1%** ViT-1D: **77.8%**
The MLP is slightly better overall; the ViT's main added value here is interpretability, not accuracy.

What did the supervised Transformer learn?

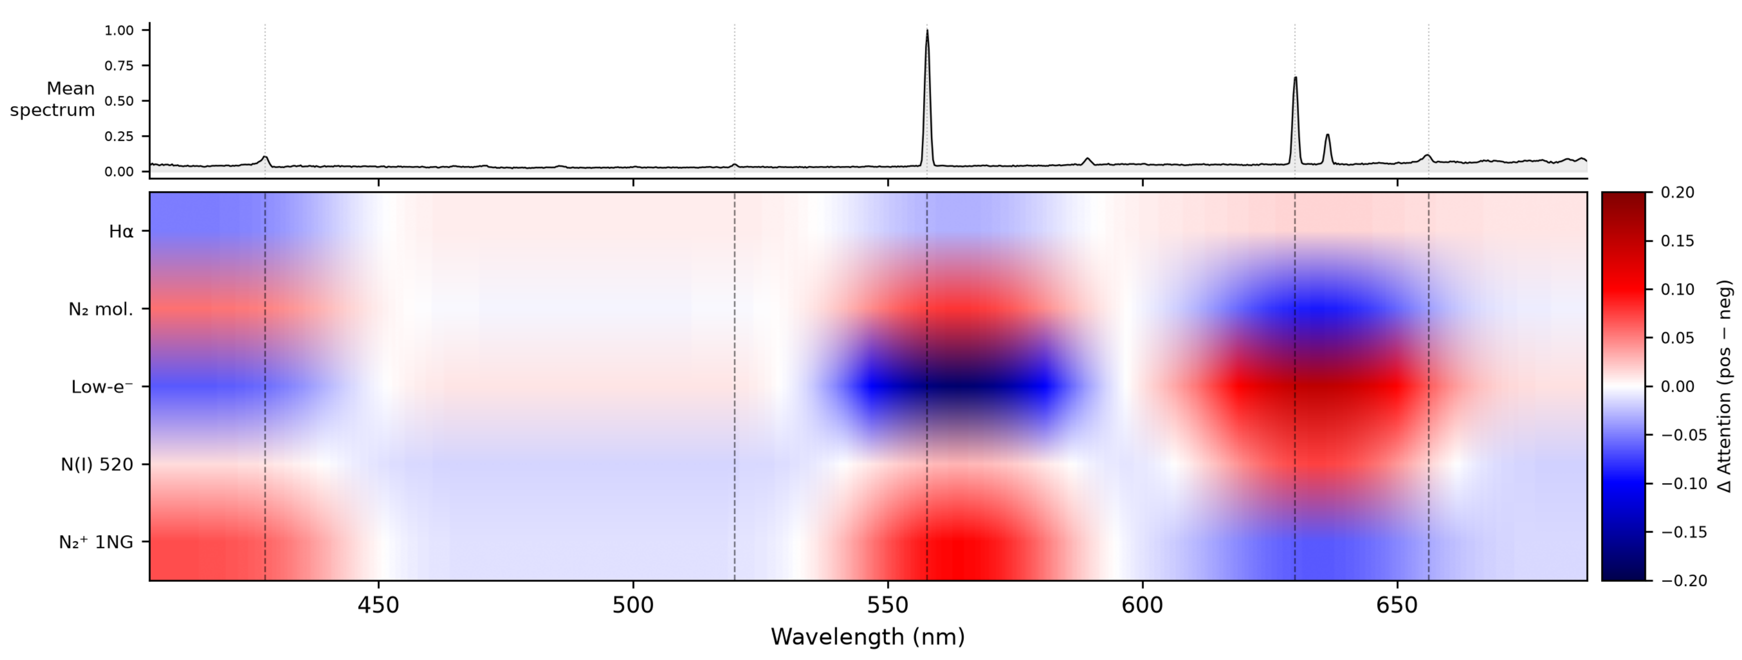

Figure: Differential attention: red regions receive more attention when a class is present; dashed lines mark known emissions.

- The model independently focuses on physically meaningful regions: N_2^+ near 427.8 nm, Hα near 656.3 nm, and broad red-OI information above 630 nm.
- **Important caveat for N(I) 520 nm:** the 520-nm line itself is not clearly attended. The decision relies mainly on co-occurring red-OI emission near 630/636 nm, so this class still requires investigation.

Published proof of concept: Le Lain et al. (2026), multi-label classification of ground-based auroral spectra with 1D Vision Transformers.
Scan for paper, code and supplementary material.

2. Unsupervised exploration of the full ASIS archive

- The 719 manually labelled spectra represent only about 0.2% of the ASIS archive. Supervised classification can only identify regimes that have already been defined by an expert.
- We therefore ask the reverse question: without using any labels, does the full ASIS archive naturally organise into recurrent spectral states?
- The analysis uses 328 281 spectra acquired over 220 nights between October 2023 and January 2026.
- For each spectrum, a compact set of quantities describing the main auroral emissions, continuum properties and spectral ratios is extracted. Density-based clustering is then used to identify recurrent groups without imposing the number of classes in advance.

328 281 spectra no labels → **Spectral diagnostics** lines, continuum, ratios → **Compact representation** of each spectrum → **Density-based clustering** groups emerge from the data

Without any expert labels or predefined number of classes, the full archive separates into six recurrent spectral states

Six recurrent spectral states emerge from the archive

The cluster names are assigned after the unsupervised analysis, according to their dominant observational characteristics.

- **C1** Bright Aurora
- **C2** Red-enhanced Aurora
- **C3** Diffuse-like Aurora
- **C4** Faint Aurora
- **C5** Bright Sky Background
- **C6** Scattered Light (twilight / Moon)

C1–C4 predominantly correspond to auroral observations, while **C5–C6** correspond to non-auroral observing conditions.
C5 and C6 together represent about 26% of the archive, showing that the same analysis also provides an efficient first-pass archive-cleaning filter.

Do the four auroral clusters carry physical information?

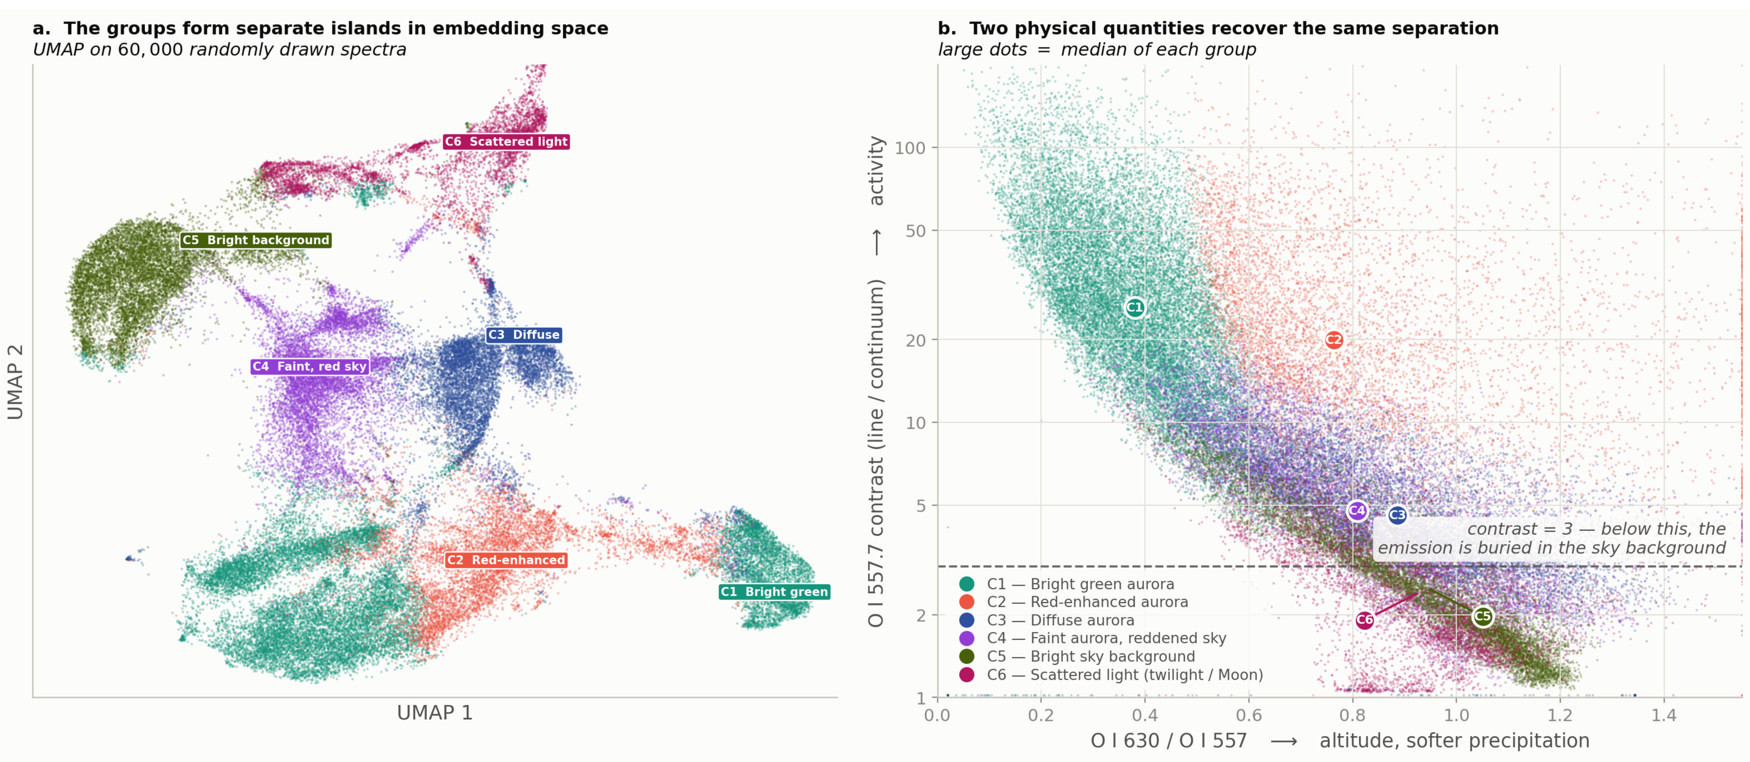

Figure: **Left:** organisation of the spectra in the low-dimensional representation (UMAP, 60 000 spectra). **Right:** the same clusters shown using two classical auroral diagnostics: the O I 630.0/557.7 nm ratio and the 557.7 nm line contrast. Large symbols indicate cluster medians.

The four auroral clusters show a clear organisation along two measurable spectral quantities:

- the **O I 630.0/557.7 nm ratio**, related to emission altitude and precipitation hardness;
- the **557.7 nm line contrast**, related to the strength of the auroral emission.

These quantities were not used as target labels, yet the clusters show a systematic ordering in this physical space.

The clustering is therefore not arbitrary: the spectral groups follow a strong physical trend.

However, the interpretation of individual clusters is not yet unique. The auroral groups appear to sample a continuum of spectral states rather than four sharply separated precipitation regimes.

Key results & outlook

- **Supervised proof of concept:** automated classification reaches about 80% class-averaged Average Precision, while the Transformer provides wavelength-resolved information on which spectral regions contribute to its decisions.
- **Unsupervised exploration:** more than 328 000 spectra organise into six recurrent states without expert labels. Four correspond predominantly to auroral observations and follow established physical diagnostics; two identify non-auroral sky conditions.
- **Toward SSL:** the existence of this physical organisation motivates learning directly from the complete spectra rather than from a predefined set of emission-line descriptors.
- **Expert in the loop:** additional annotations will target poorly represented or ambiguous regimes. An independent line-SNR module will provide quantitative evidence for diagnostic emissions such as Hα, N_2 471 nm and N(I) 520 nm alongside the AI classification.
- **Transfer to other instruments:** the methods developed here provide a general framework that can be applied to spectral datasets from other ground-based instruments and networks, including MISS, T-ReX and NORVIS.