

An Improved Artificial Neural Network CO Retrieval for IASI L2 Processor

Doc.No. : EUM/MET/TEN/09/0232
Issue : v1A
Date : 29 October 2010

EUMETSAT
Am Kavalleriesand 31, D-64295 Darmstadt, Germany
Tel: +49 6151 807-7
Fax: +49 6151 807 555
<http://www.eumetsat.int>

Document Change Record

<i>Issue / Revision</i>	<i>Date</i>	<i>DCN. No</i>	<i>Changed Pages / Paragraphs</i>
v1, v1A	29 Oct 2010		First release.

Table of Contents

1	Introduction	4
1.1	Purpose and Scope	4
1.2	Document Structure.....	4
1.3	Reference Documents.....	4
1.4	Acronyms.....	5
2	Background	6
2.1	Initial Configuration	6
2.2	First Upgrade after External Validation Feed-back	6
3	Day-2 Modifications and Improvements.....	8
3.1	Input/Output Normalisation	8
3.2	Geometry	8
3.3	Surface Elevation	9
3.4	Spectral Information	10
3.5	Teaching Database	12
3.6	Network Structure, Dimension and Training.....	13
3.7	Performances	13
4	Initial Validation with Satellite Data	15
4.1	Data Description	15
4.1.1	IASI Retrievals	15
4.1.2	MOPITT Products	15
4.2	Sensitivity to the Scanning Geometry	16
4.3	Sensitivity to the Surface Temperature	16
4.4	Spatial and Diurnal Coherence.....	16
4.5	Comparison with MOPITT	17
5	IASI Interpixel Differences and Extended Validation with Satellite Data	19
5.1	The Interpixel Differences in the L2 CO and L1C Products	19
5.2	Validation with Satellite Data	20
5.2.1	Self-consistency Checks on IASI Products	20
5.2.2	Comparison with MOPITT CO Products.....	20
6	Conclusions	21
6.1	Summary and Recommendation	21
6.2	Practical Implementation into Operational PPF.....	21
6.3	Discussion and Perspectives.....	21
Appendix A	Training and Validation Results	23

1 INTRODUCTION

1.1 Purpose and Scope

The purpose of this document is to describe and propose an improved carbon monoxide (CO) retrieval based on artificial neural networks (ANN) for implementation in the operational IASI Level 2 (L2) Product Processing Facility (PPF).

1.2 Document Structure

Section 1 is this introduction.

Section 2 gives a historical overview of the scheme being currently operated and the reasons which motivated the recent investigations for improving on this basis.

Section 3 describes the nature of the changes and summarises their theoretical performances.

Section 4 presents some first validation results.

Section 5 extends the validation with satellite data, and examines the interpixel differences discovered during the investigations.

Section 6 is a conclusion giving recommendation for implementation, and lists the items (e.g. documents, code and data) which would be affected.

Finally, a more exhaustive set of results is given in an annex, Appendix A, in the form of graphs and maps.

1.3 Reference Documents

- RD 1** Watts, P.D., M.R. Allen, T.J. Nightingale, 1996, "Wind speed effects on sea surface emission and reflection for the Along Track Scanning Radiometer", *Journal of Atmospheric and Oceanic Technology*, 13, 126-141
- RD 2** Masuda, K., T. Takashima, Y. Takayama, 1988, "Emissivity of pure water and sea waters for the sea surface in the infrared window regions", *Remote Sensing of Environment*, 24, 313-329
- RD 3** Clerbaux, C., J. Hadji-Lazaro, S. Turquety, G. Mégie, C. Camy-Peyret, S. Payan, "IASI Trace gases retrieval algorithm", Technical Document, 2002
- RD 4** Clerbaux, C., J. Hadji-Lazaro, S. Turquety, G. Mégie, "Algorithme d'inversion gaz traces pour IASI", Technical Document, Oct. 2000
- RD 5** Flemming, J., A. Dethof, M. Suttie, "Test of IASI CO total column retrievals for assimilation in the GEMS NRT suite", 28 October 2008
- RD 6** Hultberg, T., T. August, "CO and O3 total column retrieval comparison", EUM/OPS-EPS/TEN/08/0700
- RD 7** LeCun, Y., L. Bottou, G.B. Orr and K.-R. Müller, "Efficient backprop" in "Neural Networks: tricks of the trade", Springer, 1998
- RD 8** Schlüssel, P., "EPS Ground Segment IASI Level 2 Product Generation Specification", EPS.SYS.SPE.990013

- RD 9** Brasseur, G.P., D.A. Hauglustaine, S. Walters, P.J. Rasch, J.-F. Müller, C. Granier and X.X. Tie, “MOZART: A global chemical transport model for ozone and related chemical tracers, Part 1. Model Description”, *Journal of Geophysical Research*, 103, 28,265–28,289, 1998
- RD 10** Emmons, L.K., D.P. Edwards, M.N. Deeter, J.C. Gille, T. Campos, P. Nédélec, P. Novelli and G. Sachse, “Measurements of Pollution In The Troposphere (MOPITT) validation through 2006”, *Atmos. Chem. Phys. Discuss.*, 8, 18091–18109, 2008
- RD 11** Hadji-Lazaro, J., C. Clerbaux and S. Thiria, “An inversion algorithm using neural networks to retrieve atmospheric CO total columns from high-resolution nadir radiances”, *JGR*, 104, D19, pp.23841–23854, 1999
- RD 12** Turquety, S., J. Hadji-Lazaro and C. Clerbaux, “First satellite ozone distributions retrieved from nadir high-resolution infrared spectra”, *GRL*, vol.29, n°24, 2198, pp.51-1 – 51-4, 2002
- RD 13** de Laat, A.T.J., J. Lelieveld, G.J. Roelofs, R.R. Dickerson and J.M. Lobert (2001), “Source analysis of carbon monoxide pollution during INDOEX 1999”, *J. Geophys. Res.*, 106(D22), 28, 481–28, 495
- RD 14** Pochanart, P., O. Wild and H. Akimoto, “Air Pollution Import to and Export from East Asia” in “Handbook of Environmental Chemistry”, Springer Berlin / Heidelberg, Volume 4G/2004, pp.99-130
- RD 15** “IASI L0/L1 NRT monitoring”,
<http://oiswww.eumetsat.org/epsreports/html/index.php?instrument=IASI&year=2010>

1.4 Acronyms

ANN	Artificial Neural Network
ECMWF	European Centre for Medium-Range Weather Forecasts
EFOV	Elementary Field Of View
EPS	EUMETSAT Polar System
EOF	Empirical Orthogonal Function
EURD	End User Requirements Document
IASI	Infrared Atmospheric Sounding Interferometer
IFOV	Instantaneous Field Of View
LST	Land Surface Temperature
L2	Level 2
MLP	Multi-Layer Perceptron
MODIS	Moderate Resolution Imaging Spectroradiometer
MOPITT	Measurements of Pollution In The Troposphere
OEM	Optimal Estimation Method
PGS	Product Generation Specification
PPF	Product Processing Facility
SA	Service d’Aéronomie
SST	Sea Surface Temperature
TOA	Top of the atmosphere
WV	Water Vapour

2 BACKGROUND

2.1 Initial Configuration

Such ANN retrievals are basically statistical non-linear regressions between a targeted output, here the CO total column, and a collection of inputs meant to capture all the required physics it relates to. The regression is performed with a set of coefficients triggering the connections between the inputs and outputs, namely the weights and biases. The inputs and outputs usually consist of normalised parameters in the scope of the neural network itself. The weights and biases governing the retrievals are adjusted in a so-called training phase, where teaching pairs of input and output are iteratively presented to the network.

The initial configuration implemented in the PPF was defined after the specifications delivered by “Service d’Aéronomie” (SA) [RD 3, RD 4], which form part of the IASI L2 Product Generation Specification (PGS). Beyond the CO itself, the ANN approach is used to retrieve other trace gases: O₃, CO₂, CH₄ and N₂O. The first SA delivery covered networks and their associated coefficients aiming at the total columns of CO, CH₄ and O₃ as well as the partial O₃ column below 12 km. Dedicated nets were then built and trained in-house on the same basis to complete the chemistry products: the two remaining partial O₃ columns (ground to 6 km and ground to 16 km) as well as the CO₂ and N₂O total columns which are planned in the IASI L2 data.

A late delivery addressing the O₃ partial columns < 6 and < 16 km came after the first International IASI Conference. The work performed and shared by SA in the area of ANN retrievals was actually dedicated to the nadir views [RD 11, RD 12]. It was therefore eventually decided to re-train at EUMETSAT also the CO and O₃ networks to allow for retrievals at all scan angles as strong dependences on the geometry were foreseen otherwise, which were confirmed recently (see Figure 1). The architecture of the nets and the nature of inputs were kept identical to the original specifications provided by the experts. The first re-training was done with help of a database describing a wide variety of atmospheres and their associated simulated IASI spectra. Those were computed for various geographical and geometrical configurations with the forward model RTIASI-5.3.

These settings went into a pre-operational mode and the first chemistry products were released in Spring 2008 to let external partners support their validation.

2.2 First Upgrade after External Validation Feed-back

In a memo issued in Autumn 2008 [RD 5] closing an intercomparison exercise between EUMETSAT, SA, MOPITT CO products and CO models, ECMWF commented on the poor quality of the CO information provided in EUMETSAT IASI L2 packets. After a mishandling of the cloud information had been clarified, the exercise was repeated again and useful statistics derived. Although the large-scale patterns were correctly captured, a strong positive bias could be characterised in the CO columnar amounts.

The reasons for this were investigated and essentially found in inappropriate radiance corrections. Before entering any retrieval functions of the IASI L2 processing chain, the measured radiances are indeed tuned to account for the biases introduced by the forward model with respect to the measurements in the iterative retrieval and in the theoretically derived coefficients configuring the PPF. Those tuning parameters are obtained from a set of pairs of spectra: the true measurements on the one hand and synthetic spectra on the other. The latter are computed with a forward model – here, RTIASI – and some representation of the true atmospheric state mainly based on ECMWF analyses data. In the absence of accurate sources for the trace gas profiles, the radiance tuning derived for the corresponding spectral channels proved to be irrelevant and to be causing the strong bias reported above.

A second main issue dealt with the elevated regions. Further to some chosen radiances, the networks are fed with auxiliary inputs: the surface temperature and a coarse temperature profile at some fixed levels, some of which fall into the subsurface in the context of an elevated terrain. The initial specifications did not cover this aspect and therefore only ocean and low-land cases were retained in the first artificial network trainings. A post-processing correction accounting for the shortened atmospheres was patched and tested after the reception of ECMWF's memo. Together with the removal of the radiance tuning and a clarification by SA of what the input surface temperature was meant to be, it formed the essence of the IASI L2 PPF upgrade that became operational by the end of January 2009.

The release followed the conclusion of an intercomparison exercise with internal and external products, documented in [RD 6], which confirmed sensible improvements. Various retrievals from SA and EUMETSAT based on neural networks and optimal estimation methods (OEM), covering the last ten days of August 2008, were compared to MOPITT CO L3 daily products. Despite some remaining issues such as an angular dependence (see Figure 2), a smaller but persistent positive bias or occasional unrealistic columns over deserts or ice covers, the neural network retrieval was found to be more than a complement and actually a potentially good alternative to the OEM – at least in its current status – in terms of computation time, range of total columns covered and even accuracy in the case of ozone. This motivated the subsequent investigations and developments, allowing for possible changes to the initial specifications, to improve more on this ANN scheme.

3 DAY-2 MODIFICATIONS AND IMPROVEMENTS

3.1 Input/Output Normalisation

As often advised in the literature [RD 7], artificial nets are advantageously taught with pairs of inputs and outputs centred and normalised between -1 and 1. Such an individual preparation enables the maximum benefit to be obtained from each input element's dynamic and allows faster as well as more accurate learning. Normalised inputs (E_i), for instance, are typically computed as follows:

$$E_i = 0.9 \times \frac{\text{Input}(i) - C_1}{C_2} \quad (1)$$

Associating a dedicated pair of normalisation coefficients (C_1 ; C_2) to each input element then constituted the first modification to the original settings (whereby a unique pair was originally applied to the entire vector). Figure 3 illustrates how the respective ranges of the various inputs can be significantly different from each other. Furthermore, each E_i statistically now scales from -1 to 1, while it initially lay between 0 and 1. (C_1 ; C_2) have been redefined as follows:

$$C_1(i) = (\text{percentile}(\text{element}(i),97.5\%) + \text{percentile}(\text{element}(i),2.5\%)) / 2 \quad (2)$$

$$C_2(i) = (\text{percentile}(\text{element}(i),97.5\%) - \text{percentile}(\text{element}(i),2.5\%)) / 2 \quad (3)$$

Similarly, the outputs are prepared and normalised for training in the same way, such that the raw output of the neural net has to be rescaled to the expected units before entering the final product. In the initial scheme, outputs were simply normalised by dividing the columns by their order of magnitude, e.g. 10^{18} molecules/cm² in the case of CO.

3.2 Geometry

In this area, the SA approach to account for the scanning angle consists of some post-processing corrections applied to net outputs. A brief description and some empirical coefficients were communicated in an email to EUMETSAT towards the end of 2008, without further references. They mostly cover CO but do not appear to have entirely solved the undesirable effect, as can be seen in Figure 1 and Figure 2. The angular dependence in SA-NN is of the same order as observed in the current (not corrected for) EUMETSAT CO products.

The problem was solved recently by adding an extra predictor to the input vector, namely the satellite zenith angle in the form of its secant. New trainings of a single net on teaching patterns including all viewing angles with this additional input proved very successful. This is well illustrated in Figure 4 where mean retrieved CO total columns computed for each scan step, from about 100 orbits covering the last week of August 2008, show a rather flat profile. A positive side effect of this was also a more efficient teaching process, with a significant decrease of the training error.

The solar zenith angle was tested as well but did not appear to be triggering the final network output. This is well illustrated in Figure 7 which displays the weights of the connections

linking the various inputs to the inner part – the first hidden layer – of a neural network after training. The coefficients weighting the contribution of the solar zenith angle oscillate closely around zero and are two to three orders of magnitude smaller than the weights applied to the other inputs: radiance, temperature and satellite zenith angle. This is also reflected in the Jacobians of the system, of which an example is given in Figure 8, where the same considerations apply: derivatives of the output with respect to the solar zenith angle are close to zero and are two to three orders smaller than the derivatives to the other inputs.

3.3 Surface Elevation

As recalled in the first section, the net currently running operationally has not been specifically trained with elevated cases, with the result that the shortened columns are currently handled with an empirical work-around, which qualitatively has proved suitable in a first instance.

Training and retrieving the total column when the surface pressure falls below 980 hPa requires handling properly the portion of the input temperature profile which lies below the surface. In the absence of particular guidance from the originators in this area, two approaches were tested. In the first solution, the input temperature profiles were defined on a variable grid whose levels were dynamically defined according to the surface pressure, such that no levels ended up in the subsurface. The other alternative, which was eventually retained, is based on a fixed grid – as initially specified – where any subsurface component extends an isothermal profile from the first level above the surface to the lowest pressure level foreseen. Strictly, this extension should be done with the brightness temperature of the surface. In such a configuration, if it was real, the radiative transfer equation would impose the restriction that the contributions from the lowest part could not be distinguished and would then be masked. This is what the neural network is expected to capture and reproduce. It however requires an accurate knowledge of the surface temperature together with the surface emissivity, whose retrieval has not been validated over land yet. As a work-around, the first atmospheric temperature is then used here as proxy information in the proposed implementation. The selected pressure levels are listed in Table 1.

To support these particular retrievals, an additional predictor was introduced to the input vector: the surface pressure. Further trainings were repeated and included a collection of various continental elevated cases. Verifications with real data show a good correlation between lower surface pressure and reduced CO columns: the amounts are consistently smaller over mountains than in the surrounding low lands. In Figure 16, the biggest formations like the Himalayas and Andes are qualitatively well captured, as are the Alps, despite their smaller scale. As with the satellite zenith angle, Figure 7 and Figure 8 illustrate the importance of this predictor for the final retrieval.

id	level	hPa
0	87	978.982
1	85	899.686
2	83	826.576
3	82	792.184
4	80	727.436
5	78	667.708

id	level	hPa
6	76	610.600
7	75	587.638
8	73	543.053
9	70	478.540
10	68	436.950
11	62	321.500

id	level	hPa
12	56	222.940
13	51	155.428
14	45	93.2342
15	28	10.3700
16	18	1.36116
17	10	0.222228

Table 1: Level selection on RTIASI grid for the input temperature profile to the ANN CO retrieval

3.4 Spectral Information

This is the key information driving the retrievals and the channels involved split into two categories: a selection in the various CO absorption lines and a smaller set of so-called baseline channels where less absorption occurs and from which the background radiance is inferred. This concept was introduced to isolate the contribution of the CO to the extinction of the upwelling radiation. As detailed in equations 254 to 256 of the IASI L2 PGS [RD 8], radiances in the trace gas lines are basically subtracted from the baseline to form the spectral component of the input vector and enter the neural network. This was the unique explicit spectral information specified originally [RD 1, RD 2].

The basic transfer equation for radiation traversing a thin layer of a gas can be basically written as:

$$\frac{dI_v}{dS} = J_v - K_v I_v \quad (4)$$

where I_v is the radiance at the wavelength v , S is the path through the layer, J_v expresses the source function and K_v the absorption coefficient of the gas at that particular wavelength v . From Equation 4 it follows that the sought column cannot be retrieved only from the radiance extinction (dI_v) at the top of the atmosphere (TOA) and requires additional information about the upwelling radiance below the absorbing layer.

In the original specifications, the only proxy information for this was carried by a pseudo-surface temperature parsed as an auxiliary input to the net. It was computed with the inverse Planck function of the radiances measured at the so-called baseline channels and a surface emissivity artificially fixed to 0.975. However, a few absorptions, mostly due to water vapour but also to N_2O , affect these channels and decouple the resulting radiation from the pure surface emission (see Figure 9). The CO retrieval in its original form proved to suffer from these approximations over many surfaces, especially over some barren land and ice covers.

To circumvent this, we propose to replace this pseudo-surface temperature with the raw radiances picked at some so-called baseline channels, which describe the background radiance before it is attenuated by CO. The surface temperature is also kept as an explicit parameter such that the emissivity and potential interference from other gases are assumed now to be implicitly covered by the combination of the inputs and expected to be statistically captured during the network learning phase. Some tests however showed that the surface emissivity further helps the CO retrieval if explicitly added to the input vector. When this IASI product has been fully validated, this parameter could become a valuable additional descriptor of the physics fed into the neural net. Eventually, the collection of temperatures sampling the vertical profile completes this information by describing the source function in Equation 4.

The channels originally proposed by SA are listed in Table 2 and are represented in their spectral context in Figure 5. Some of them are strongly affected by the water vapour (WV) as can be seen in Figure 10 which plots the correlations between the intensity of a line and CO and WV concentrations at different levels. The updated networks were able to recognise this and gave less weight to those channels, which is well illustrated by channels 3 and 4 in Figure 7 and their corresponding Jacobians in Figure 8 for instance.

id	chan	cm ⁻¹
0	5866	2111.25
1	5867	2111.50
2	5868	2111.75
3	5869	2112.00
4	6022	2150.25
5	6023	2150.50
6	6024	2150.75
7	6037	2154.00
8	6038	2154.25
9	6039	2154.50

id	chan	cm ⁻¹
10	6052	2157.75
11	6053	2158.00
12	6054	2158.25
13	6055	2158.50
14	6056	2158.75
15	6081	2165.00
16	6082	2165.25
17	6083	2165.50
18	6084	2165.75
19	6085	2166.00

id	chan	cm ⁻¹
20	6096	2168.75
21	6097	2169.00
22	6098	2169.25
23	6099	2169.50
24	6111	2172.50
25	6112	2172.75
26	6113	2173.00
27	6114	2173.25
28	6126	2176.25
29	6127	2176.50

Table 2: Initial channel selection. Absorption lines are listed above and baseline channels below.

0	5955	2133.50
---	------	---------

1	5993	2143.00
---	------	---------

2	6021	2150.00
---	------	---------

This selection was eventually modified in a final step (see Table 4 and Figure 6) to provide to the network some explicit information about the interfering species, namely water vapour and N₂O here, and thus support the discrimination of the CO contribution. It comes as a selection of extra baseline channels free of CO features with independent or combined H₂O and N₂O lines. Further CO channels could also be added with benefit, and the respective improvements in terms of learning error are summarised in the following table.

Config. Gas	Original SA settings	Day-2 improvements + initial channel selection	Day-2 improvements + modified channel selection
CO [e18 molec/cm ²]	0.370 ¹	0.270 ²	0.235 ²

¹ ground elevations excluded (P_s < 980hPa)

² all elevations are included in the training sets

Table 3: Evolution of the ANN CO training error on TCE with the successive algorithmic and configuration improvements

id	chan	cm ⁻¹
0	5865	2111.00
1	5866	2111.25
2	5867	2111.50
3	5868	2111.75
4	5869	2112.00
5	5870	2112.25
6	6022	2150.25
7	6023	2150.50
8	6024	2150.75
9	6025	2151.00
10	6026	2151.25
11	6037	2154.00
12	6038	2154.25
13	6039	2154.50
14	6040	2154.75
15	6041	2155.00
16	6052	2157.75
17	6053	2158.00

id	chan	cm ⁻¹
18	6054	2158.25
19	6055	2158.50
20	6056	2158.75
21	6081	2165.00
22	6082	2165.25
23	6083	2165.50
24	6084	2165.75
25	6085	2166.00
26	6095	2168.50
27	6096	2168.75
28	6097	2169.00
29	6098	2169.25
30	6099	2169.50
31	6100	2169.75
32	6101	2170.00
33	6109	2172.00
34	6110	2172.25
35	6111	2172.50

id	chan	cm ⁻¹
36	6112	2172.75
37	6113	2173.00
38	6114	2173.25
39	6123	2175.50
40	6124	2175.75
41	6125	2176.00
42	6126	2176.25
43	6127	2176.50
44	6128	2176.75
45	6129	2177.00
46	6138	2179.25
47	6139	2179.50
48	6140	2179.75
49	6141	2180.00
50	6142	2180.25
51	6143	2180.50

Table 4: Upgraded channels selection. CO absorption lines are listed above and baseline channels below.

id	chan	cm ⁻¹
0	5985	2141.00
1	5986	2141.25
2	5987	2141.50

id	chan	cm ⁻¹
3	5988	2141.75
4	5989	2142.00
5	5990	2142.25
6	5991	2142.50

id	chan	cm ⁻¹
7	5992	2142.75
8	5993	2143.00
9	5994	2143.25

10	5995	2143.50
11	5996	2143.75
12	5997	2144.00
13	5998	2144.25
14	5999	2144.50
15	6000	2144.75

16	6116	2173.75
17	6117	2174.00
id	chan	cm⁻¹
18	6118	2174.25
19	6119	2174.50
20	6120	2174.75

21	6131	2177.50
22	6132	2177.75
23	6133	2178.00
24	6134	2178.25
25	6135	2178.50

3.5 Teaching Database

The training database approximately contains 200,000 patterns made up of atmospheric state vectors and their associated synthetic IASI spectra computed with RTIASI-5.3. The atmospheric temperature, humidity surface pressure and wind components are based on the Le Chevallier climatological database. Neither clouds nor aerosols were included such that the networks learnt pure clear cases only, which are also subsequently their only domain of validity.

Some missing information, e.g. the trace gas profiles, was forged to cover the whole range of expected situations with random variations around standard profiles. In the case of the CO, the vertical distributions are based on 43 original profiles sampled from the MOZART 3D chemical transport model calculations [RD 9], obtained from D. Cunnold's runs in 2001. These include temperature profiles, so that for each of the synthetic cases the CO profile was chosen which had the closest surface air temperature. To generate a realistic and continuous set of scenarios for CO, this selected profile was subsequently randomly either left unchanged or varied by adding up to half of the variability (max - min) of the mixing ratio in the basic 43 modelled vertical distributions.

Figure 29 to Figure 32 present the statistics and distribution of the resulting CO profiles and how they compare with MOPITT daytime retrievals during the whole year 2008. It can be seen that the range of teaching CO abundances generally cover the real situations well, with the exception of a few extreme events of very high concentration which the training set does not include yet. When compared to MOPITT ones, the synthetic profiles are also possibly too skewed, showing a slight relative excess towards the surface. However, MOPITT information was used here without accounting for its vertical sensitivity, which falls at the lower levels and could explain the more vertical shape of the satellite products. The correlation between lower and upper layers is higher in the training base than in the MOPITT data, respectively amounting to 0.7 and 0.4 approximately, whereas the correlations between adjacent layers are similarly high in both data sets.

Figure 30 shows the geographical distribution of the teaching cases, which cover sea as well as land situations. Different scan angles were randomly assigned to those cases in order to equally cover all of the instrument viewing geometries. The synthetic IASI spectra at the top of the atmosphere were successively computed with and without solar contributions to simulate day and night situations. In the case of daytime, the solar angles were randomly defined to equally represent the Sun elevations.

Elevated areas are by essence included in the training set which also comprises a wide range of surface types. Over water, the surface emissivity was computed after Masuda's model (RD 2) and its extension by Watts (RD 1). Ground emissivities were derived from the MODIS UCSB emissivity library, which was downloaded from

www.icesb.ucsb.edu/modis/EMIS/html/em.html. Based on the spectra of pure surfaces, composite surfaces have been generated with random contributions from up to three different types, excluding however some combinations like snow/ice at tropical temperatures. The emissivities have been further varied by random numbers, reducing the spectral emissivities by 10% on average (total range 0 to 20%). The pure sand ground type, with unique features around 8-9 μm , was not retained for this initial simulation and will be included in coming upgrades of the neural network configuration coefficients.

3.6 Network Structure, Dimension and Training

The ANN used here are multi-layer perceptrons (MLP) made of an input and an output layer connected through two hidden layers. This follows the initial design and so do the activation functions retained for the hidden (*tanh*) and output (*identity*) layers. It was however found appropriate to increase the size of the hidden layers from 8 to 48 neurons. Adding even further neurons helped the training but made the network too specific to the learning sets. The collection of inputs after the new design is summarised hereafter and the unique output is the total column of CO. The channel and pressure level selections are configurable and can be modified without changing the implementation.

Absorption lines (52 channels)	Baseline radiance (26 channels)	Surface temperature	Temperature profile (18 levels)	Satellite zenith angle (secant)	Surface pressure
-----------------------------------	---------------------------------------	------------------------	---------------------------------------	---------------------------------------	---------------------

Table 5: Input vector to the ANN CO retrieval

The teaching database is split into a training and a control set. To avoid overtraining, i.e. the net becoming too specific to the teaching patterns and losing its generalisation ability, one monitors the retrieval error of the control set and stops the learning when it starts diverging from the training error. This is typically achieved within approximately a couple of days with a full CPU on a machine such as TCProtos.

The teaching patterns are randomly selected to get a distribution of target CO columns as flat as possible (see Figure 11) in order to give all potential abundances the same chance to be correctly trained for and therefore retrieved later on. No restrictions apply to the elevation or the geographical location (polar situations are included) and it was ensured that no artificial correlations between the various input elements were introduced.

In order to prepare the nets for real conditions in operations, some noise is added to the input radiances, based on the instrument noise characteristics. Similarly, the auxiliary temperature inputs are degraded with some Gaussian errors amounting to 1.25 K in the case of the teaching vertical profiles, and respectively 0.4 and 2 K for sea and land surface temperature (SST and LST). Noise- and error-free trainings were also tested with the sole goal of verifying the overall approach. They obviously gave much better theoretical results with the training base (see Figure 26 to Figure 28) and lead to unstable retrievals if applied to real measurements, as expected.

3.7 Performances

The theoretical performance of such artificial neural networks can be characterised with the synthetic data set that served in the training process by comparing the retrieved columns to the targets. Results for the proposed network are detailed in Figure 12 to Figure 14. The

correlation between retrieved and target columns is 0.99 and the linear relation is very close to unity. Due to the non-linear nature of the MLP, the errors are not statistically Gaussian-like, as can be seen in Figure 14. The rms of the absolute errors typically ranges between 0.2 and 0.24×10^{18} molec/cm² and appear to be quite independent of the column density itself. As a result, the relative errors are strengthened for the thinnest columns. They however generally remain well below the thresholds specified in the EPS End User Requirements Document (EURD), set to 20%. Excluding the 7% faintest columns ($< 0.7 \times 10^{18}$ molec/cm²) from the error characterisation does not change the absolute error numbers and gives an estimation of the relative error between 7 and 11%.

The distribution of the errors is also non-symmetrical such that, although the overall average is close to the origin, the main mode presents a small positive bias of approximately 1 to 3%. This is well visible in Figure 14. After investigations, this was attributed to the non-linear response to the inputs with respect to the target columns and in particular to the noise perturbing the radiances. Indeed, noise- and error-free trainings induced a rather more centred and symmetrical error distribution (see Figure 27 and Figure 28). These are pure theoretical figures though, which were only computed to assess the impact of the noise. They are not applicable to operations as in practice measurements and auxiliary temperatures are always affected by some errors.

In practice, some coefficients can be manually adjusted to account for this asymmetry with the proposed network and to remove the bias for the most commonly-occurring CO columns (1 to 4×10^{18} molec/cm²).

4 INITIAL VALIDATION WITH SATELLITE DATA

4.1 Data Description

4.1.1 IASI Retrievals

The trained network was applied to series of IASI measurements covering the last week of August 2008, and only clear IASI IFOVs as identified by the IASI L2 PPF “Cloud test A” were processed. The same network was applied twice, first with input temperatures coming from an empirical orthogonal function (EOF) stand-alone retrieval and then from ECMWF analyses products. They are respectively labelled 066 and 067 in the figures documented in Appendix A. In both cases, the surface temperature came from the IASI L2 EOF regression retrieval, which is believed to be generally the most accurate option. Ingesting ECMWF temperatures, especially the analyses, is excluded in the operational framework which currently foresees EOF in that respect. As described in §3.6, some error characteristics of the input temperatures retrieved with the EOF method were represented during the training. The retrieval 067 was then solely run here to distinguish the errors in the CO retrieval that would be caused by not treating these profiles and surface temperature properties exhaustively, e.g. ignoring potential systematic biases or a possible small angular dependence. The IASI CO retrievals were not tuned to remove the little positive bias described in §3.7.

4.1.2 MOPITT Products

The retrievals were compared over that period on an IFOV basis to MOPITT L3 Daily gridded (1°x1°) CO products (v3) downloaded from the NOAA WIST server (<https://wist.echo.nasa.gov/>). They result from an optimal estimation method and were handled here without their averaging kernels. Such products contain both the vertical profile and its integration into the total column, and also detail on a level basis how much the final retrieval (CO profile) is constrained by the *a priori* information.

In order not to be biased by the background used in the MOPITT product generation, an additional filter was applied to reject all retrievals constrained by more than 50% on average. This mainly concerns the elevated polar regions as can be seen in Figure 20. These maps also illustrate some regional and diurnal variations of such retrievals. Over oceans, the *a priori* information seems to be increasingly contributing to the product with latitude, while continental columns are more constrained by the background during night-time (~40-45%) than they are when the sun shines (~20-25%). This implies that the optimal estimation method has more degrees of freedom when the ground thermal contrast is high, which is mainly achieved over warm continental places under sunlight. On the other hand, as the surface temperature decreases, the system loses its sensitivity to the lower atmospheric layers (see Figure 21) and the blind portion of the profile is mostly inferred from the background information.

The detailed results of the intercomparison and statistics of the departures are displayed from Figure 15 to Figure 25.

4.2 Sensitivity to the Scanning Geometry

Statistics were computed separately for each viewing position (EFOVs) during that period and for day and night-time respectively. Averaged over the approximately 100 orbits, the mean CO column is expected to be approximately constant with scan angle. Furthermore, densities retrieved during daytime are also expected to match those retrieved at night time. The short period addressed here does not offer daylight sensing over Antarctica and conversely, no night counterparts are available for North Pole measurements. Therefore the polar caps were excluded from this particular angular analysis.

As can be seen in Figure 18, the implementation in PPF v4.3.2 (013b) shows a significant angular dependence, with increasing columns towards the nadir. The amplitude of the variations is as high as 0.3×10^{18} molec/cm² during daytime and even up to 0.4×10^{18} molec/cm² for night cases. The proposed implementation (066) shows a much flatter profile in general, with a small decrease less than 0.1×10^{18} molec/cm² at the swath edges for both day and night retrievals. These features however do not show with ECMWF analyses temperatures (retrieval 067), whose production is essentially independent of the IASI viewing geometry.

4.3 Sensitivity to the Surface Temperature

Because it determines the upwelling radiance transferred through the atmosphere, the surface temperature parameter is potentially a strong support for the ANN CO retrieval, as recalled in §3.4. A sensitivity study was therefore conducted where the skin surface temperature was artificially varied by ± 1 to 5 K around the EOF retrieved values before entering the ANN CO retrievals for the last seven days of August 2008. Results are presented in Figure 19 and illustrate the importance of the channel selection and of modelling the input temperature errors. With the initial channel selection and with the assumption at training of error-free input temperatures (052's network), the final CO columns responded to a 5 K change of the surface temperature by about 0.25×10^{18} molecules/cm² (~15%) on average and showed a sensitive dependence on the scan angle, which also applies to a lesser extent to their standard deviations. The proposed network (used in retrieval 067) is configured with the upgraded spectral information described in Table 4 and the teaching input temperatures used for its training were degraded with some errors. It proves much more robust with a variation of only 0.075×10^{18} molecules/cm² for the average CO columns in response to a surface temperature modification of 5 K. This good consistency as well as the stability with the scanning geometry described in §4.2 also applies to the standard deviations of the CO columns accumulated during the studied period.

4.4 Spatial and Diurnal Coherence

With an atmospheric lifetime of about two months, day and night-time CO retrievals are expected to generally match each other. With the exception of some desert sub-regions in Australia, Mongolia or Sahara, the day/night contrast is globally relatively small on the averaged maps (Figure 15 and Figure 16) and fluctuates around 0.05×10^{18} molec/cm² (Figure 18). This number is essentially indicative because clear land cases are less represented at night and because the exact same Earth points were not necessarily associated with retrievals from the day and night overpasses. In particular, the two main sources in

South America and Central Africa are mostly retrieved at daytime and were more discarded at night-time by the cloud detection for instance.

In the same way, low land and sea retrievals should present smooth transitions in coastal areas. This is in general achieved with only a few exceptions, for instance off the coasts of Morocco and Mauritania. In that particular case however, MOPITT reported the area as cloudy and consequently did not provide any CO measurements.

While sources and atmospheric transport will cause natural CO variations with time at a given location, certain regions of the globe, like the Central Pacific, appear to be constant on large scales, with lower amounts on average in this particular case. The spatial coherence of the ANN CO retrievals was studied over this area, whose standard deviation (mapped in Figure 17) may be interpreted as an approximate measure of the actual retrieval error. For the studied period, this ranges between 0.18 and 0.24×10^{18} molec/cm² on average, which is close to but slightly higher than the theoretical figures detailed in §3.7. On the other hand, the retrievals over Antarctica show much bigger and unlikely standard deviations of about 0.5 and up to 0.8×10^{18} molec/cm² in the operational configuration (066). The neural network fed with ECMWF temperature profiles (067) has a slightly lower variability at the South Pole. It amounts to approximately 0.25 over the elevated part but still remains as high as 0.5×10^{18} molec/cm² along the coasts. The low temperatures, the icy cover as well as thinner CO columns degrade the signal-to-noise ratio and would partly explain the instability of these retrievals together with potential cloud contamination, whose detection is less accurate at night under such polar conditions.

The known CO sources [RD 13] are qualitatively well retrieved. They essentially split into two main classes: the biomass burning below the Equator in Africa and South America on the one hand, and the agricultural fires and industrial emissions over Northern India and China on the other. The latter plumes are transported out over the Pacific Ocean beyond Japan [RD 14].

4.5 Comparison with MOPITT

The intercomparison with MOPITT products summarised hereafter is more exhaustively displayed in Figure 22 to Figure 25. It was broken down into day/night and geographical classes covering the Poles ($|\text{lat.}| > 60^\circ$), mid-latitudes ($30^\circ < |\text{lat.}| < 60^\circ$) and a wide tropical-equatorial band from 30°S to 30°N. Visually, MOPITT products present noticeable day/night contrasts over central Africa, Greenland and Asia for instance, and there appear to be fewer continental night measurements in general available from both instruments.

The proposed implementation (066) departures from MOPITT products are slightly higher than the theoretical errors, with standard deviations varying between 0.25 and 0.32×10^{18} molec/cm² in general, and up to 0.39 in the Northern Hemisphere at night. This translates into approximately 15 to 22% in relative terms, with the exception of Antarctica. Their distributions are usually not symmetrical and small tails are visible where IASI ANN CO is in excess, whose root cause was mainly found in the interpixel dependency discussed in Section 5. Although the standard deviation of the intercomparison only varies by a few percent in the various geographical and illumination classes, the bias can be very different from one class to another. It reaches 25% in Antarctica while amounting to 13% in the

Southern Hemisphere, 7% at the North Pole and almost vanishes between 30°S and 60°N. In terms of correlations, the best results are obtained during daytime.

Although giving good indications, no definitive conclusions can be drawn from this exercise in terms of absolute calibration. Indeed, as described in §4.1.2, the MOPITT retrievals are progressively constrained by the background information as the surface temperature drops. Further to this, recent studies showed that the MOPITT CO is overestimated [RD 9], which would imply that the proposed implementation could statistically also have a positive bias.

5 IASI INTERPIXEL DIFFERENCES AND EXTENDED VALIDATION WITH SATELLITE DATA

Systematic differences have been found between the four pixels, which will be presented and discussed in this section. The IASI L1C products indeed exhibit some interpixel differences around 2100 cm^{-1} which impact the IASI L2 carbon monoxide (CO) retrievals. The validation study presented in Section 4 was then repeated for pixels 3 and 4 only and extended to the full months of August and November 2008 and February 2009.

5.1 The Interpixel Differences in the L2 CO and L1C Products

Figure 33 shows a close-up view of the CO total column product over Madagascar on 21 September 2007 around 18:00 UTC. IASI acquires four spectra simultaneously, which are associated with individual fields of view (IFOVs, also called pixels in this section) forming an elementary field of view (EFOV). Some EFOVs are highlighted with black ellipses to ease the identification of the respective pixels. Pixel 2 is located to the upper right of an EFOV in this scene and appears to be strongly biased in comparison to the others.

The CO total column means and standard deviations were computed for each IFOV separately in the entire orbit from which Figure 33 was extracted and are displayed on that same figure. It turns out that the retrievals with pixel 2 are about 20% higher than for pixels 3 and 4, with a significantly higher standard deviation. Given the number of pixels involved (approximately 15,000 in total) and their global distribution in latitudes, we can expect that the atmospheric situations sensed by each IFOV were statistically comparable. Also considering that the very same algorithm was applied to all IFOVs, such a pixel dependency was attributed to differences in the L1C data.

This assumption was confirmed by inspection of the L1C radiances and is illustrated by Figure 34 and Figure 35. The first plot shows the mean radiance departures for the scene presented in Figure 33 between pixels (1,2,4) and pixel 3 at channels involved in the CO retrieval. Whilst the radiances of pixels 3 and 4 are in good agreement, with no bias and a dispersion smaller than $5 \times 10^{-8}\text{ W/m}^2/\text{sr/m}^{-1}$, the radiances of pixels 1 and 2 present a dispersion about twice as high as well as some large biases from pixel 3 (up to $10^{-7}\text{ W/m}^2/\text{sr/m}^{-1}$). Figure 35 shows similar statistics computed at all IASI channels with approximately 700,000 clear sky measurements over oceans at night, accumulated during more than 2.5 years of IASI L0/L1 NRT monitoring [RD 15]. Pixels 1 (black), 3 (red) and 4 (green) are compared to pixel 2 in brightness temperature, the difference being computed for a reference temperature of 280 K. Some interpixel differences can be as high as 0.35 K between 2000 and 2250 cm^{-1} where the CO retrieval is sensitive. Here again, pixels 3 and 4 present similar behaviours which motivated the validation of the CO total columns for those IFOVs separately.

5.2 Validation with Satellite Data

5.2.1 Self-consistency Checks on IASI Products

As can be seen in the monthly means and standard deviations presented in Figure 39 and Figure 40, the conclusions drawn in §4.4 with respect to the day/night and land/sea continuity basically are still valid when only the pixels 3 and 4 are retained. The main patterns are still well captured: tropical biomass burning and pollution in the Northern Hemisphere. It is also interesting to note that the CO produced in the wild bush fires that raged in early February 2009 on the south-eastern coasts of Australia and which was transported out over the southern Pacific is well visible in the monthly maps (see Figure 39 and Figure 40, bottom line).

Given the lifetime of this gas, the spatial variance of the CO retrieved amounts over a week or a month can be used as a proxy measure of the retrieval error where no particular transports modified the background load. The standard deviations computed over the central Pacific in August 2008 dropped down to 0.13 to 0.20×10^{18} molec/cm² when restricted to IFOVs 3 and 4, which is lower than in the statistics applying to all four IFOVs (cf. §4.4 and see Figure 40 and Figure 36 vs. Figure 17). This is also consistent with the theoretical errors obtained with the training set (§3.7).

5.2.2 Comparison with MOPITT CO Products

The statistics for all four pixels for the last week of August 2008, discussed in §4.5, were computed again for just pixels 3 and 4. They are shown in Figure 37 (relative departures) and Figure 38 (absolute departures) which can be directly compared to Figure 24 and Figure 25 respectively. In relative (absolute) terms, the standard deviations decreased by 3% (0.05×10^{18} molec/cm²) if the IFOVs 1 and 2 are filtered out.

The intercomparison with MOPITT, performed on a pixel basis with the IASI products, has been extended to the entire months of August and November 2008. The results are summarised in Figure 41 to Figure 44. The global correlation is 0.82 for August 2008 with pixels 3 and 4 only, which is higher than the 0.75 reported with all pixels in §4.5. The departures between IASI and MOPITT CO products have also significantly decreased, with standard deviations ranging between 10 and 15% (0.22 and 0.28×10^{18} molec/cm²), whereas they varied between 15 and 22% (0.25 and 0.32×10^{18} molec/cm²) when all four IFOVs were considered. The bias between the two products is still variable with latitude, showing larger values in the Southern Hemisphere in general.

Taking into account an average error of about 10% in MOPITT products, these results indicate that such artificial neural network retrievals mostly match the accuracy defined in the EPS mission requirements.

6 CONCLUSIONS

6.1 Summary and Recommendation

The initial specifications for a neural network retrieval ran in a pre-operational mode for several months and showed some limitations, which were characterised. The main issues dealt with a strong scan angle dependency, a positive bias and failure over elevated areas as well as some specific surfaces (ice and deserts). Investigations were conducted, from which a collection of modifications were elaborated with a single architecture and were discussed in this document.

Interpixel differences in the IASI CO products were analysed and traced back to the L1c measurements. As a conclusion, the pixels 3 and 4 were retained and further validated together. Verification with training data and preliminary validation with real IASI measurements and external satellite products were performed which confirmed the improvements and the level of maturity in comparison to the current operational processing chain. In terms of timeliness, this technique is compatible with operational processing on the EPS ground segment. The validation errors with the pixels 3 and 4 match the required accuracy, such that it is recommended to implement these changes into the IASI L2 PPF for the revision v5.0.

6.2 Practical Implementation into Operational PPF

A few elements need to be modified accordingly, which are listed hereafter:

- The static auxiliary data (SAD) description (document EUM.EPS.SYS.TEN.03.034) needs to be updated to reflect the additional inputs and the individual scaling coefficients for every input and output element in COF_TRGAS;
- The COF_TRGAS itself;
- The “readaux” library;
- The function ANN_OADFG() within iterativeRetrieval.cc;
- The IASI L2 PGS, section 5.19.

6.3 Discussion and Perspectives

Although this report addresses CO only, it is believed that the modifications which are suggested here are applicable to the other gases foreseen in the IASI L2 products and that their retrievals should be improved as well. Similarly to the CO, the N₂O products should also be produced for pixels 3 and 4.

It is acknowledged that the validation exercise with external data covered a limited period. It should be extended to other dates and reference products. Further to this, validations with ground-based Fourier transform infrared (FTIR) spectrometer measurements of the Network for Detection of Atmospheric Component Change (NDACC) will be performed in order to complete the quantitative error characterisation on smaller scales.

Concerning the training set, the impact on the final retrieval of the correlation between bottom and upper levels, as well as the positive skewness of the synthetic profiles towards the surface, should be investigated. Indeed, because of the reduced sensitivity to the boundary layers in many situations, their contributions are mostly inferred from the upper part of the profile and the retrieved total column may be biased if the teaching vertical distribution is artificially exaggerated and correlated. Additionally, some extreme CO columns, well above 7×10^{18} molec/cm², were reported in recent presentations given by LATMOS-ULB, occurring in major events such as the wild fires in Greece, 2007. MOPITT products over the year 2008 also contain a minority of high density profiles which are outside the training range. Because of its non-linear response, it is expected that the neural network retrieval will be degraded under such circumstances. A last aspect to be considered is potential systematic biases between the fast radiative transfer model (here, RTIASI) and the actual measurements. In contrast to their application in the optimal estimation method set up in the operational IASI L2 PPF, the radiances ingested by the artificial neural network are not corrected for. Hence, we may expect biased retrievals like those observed and described in Section 5. Their characterisation will require extended comparisons with models, external satellite products and ground measurements.

Ultimately, current external works based on EOF (Dan Zhou) or OEM (LATMOS-ULB) methods actually aim at retrieving the CO profiles from IASI measurements. Although not more than 1.7 degrees of freedom were claimed to be present on average for that particular gas with this OEM approach, it could be investigated whether the ANN can be used to infer the vertical CO distribution at a cheaper computation-time cost.

APPENDIX A TRAINING AND VALIDATION RESULTS




Figure 1: Angular variation of the retrieved ozone total column over the Arabian Sea on 25 August 2008. Ozone abundances are shown as a function of the scan angle (x-axis), the nadir being indicated by the dash-dot line. Plain thick and dashed lines respectively represent retrievals obtained with neural nets operated at EUMETSAT (EUM-NN) and at former “Service d’Aéronomie” (SA-NN) now LATMOS. The off-nadir views were discarded in the latter processing chain.




Figure 2: Angular variation of the retrieved CO total column over the Pacific Ocean on 25 August 2008. Abundances are shown as a function of the scan angle (x-axis), the nadir being indicated by the dash-dot line. Plain thick and dashed lines respectively represent retrievals obtained with neural nets operated at EUMETSAT (EUM-NN) and at Service d’Aéronomie (SA-NN). The swath edges are discarded in the latter processing chain.

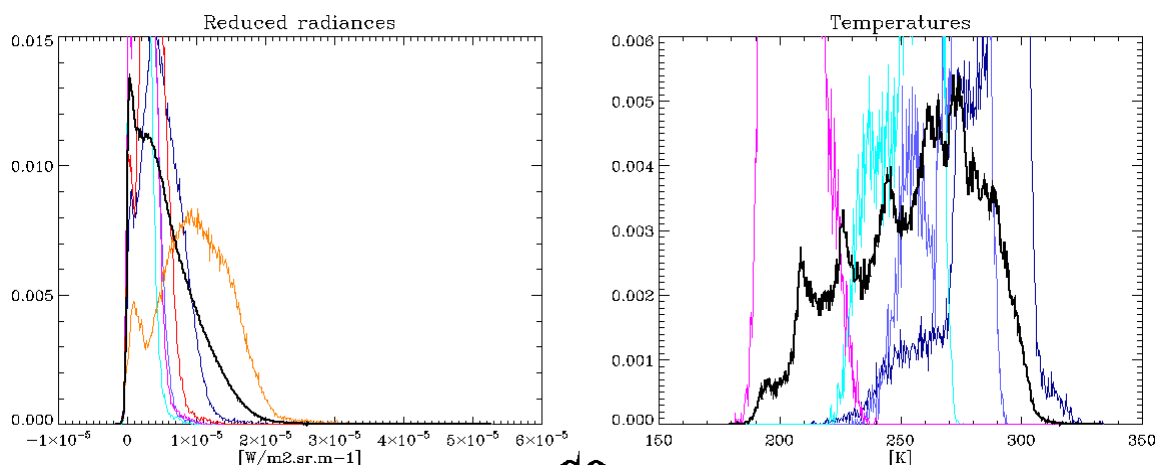


Figure 3: Distribution of some radiance (left) and temperature (right) inputs to the ANN CO retrieval. Different colours are associated with individual input parameters (channels or levels) while the thick black lines draw the overall statistics for all involved radiances and atmospheric temperatures.




Figure 4: Angular variation of the mean CO total column retrieved with IASI acquisitions in the last week of August 2008 under various ANN configurations. 013b is similar to the current operations while 047, 052 and 053 are candidates for the next PPF upgrade and use the satellite zenith angle as an explicit predictor. + and • symbols correspond to night and day times, respectively.




Figure 5: Original channel selection. Pink and black vertical dashed lines highlight the CO absorption and baseline channels involved in the ANN as originally selected by SA. Transmittances of various atmospheric components in the CO spectral region were computed for a tropical case.

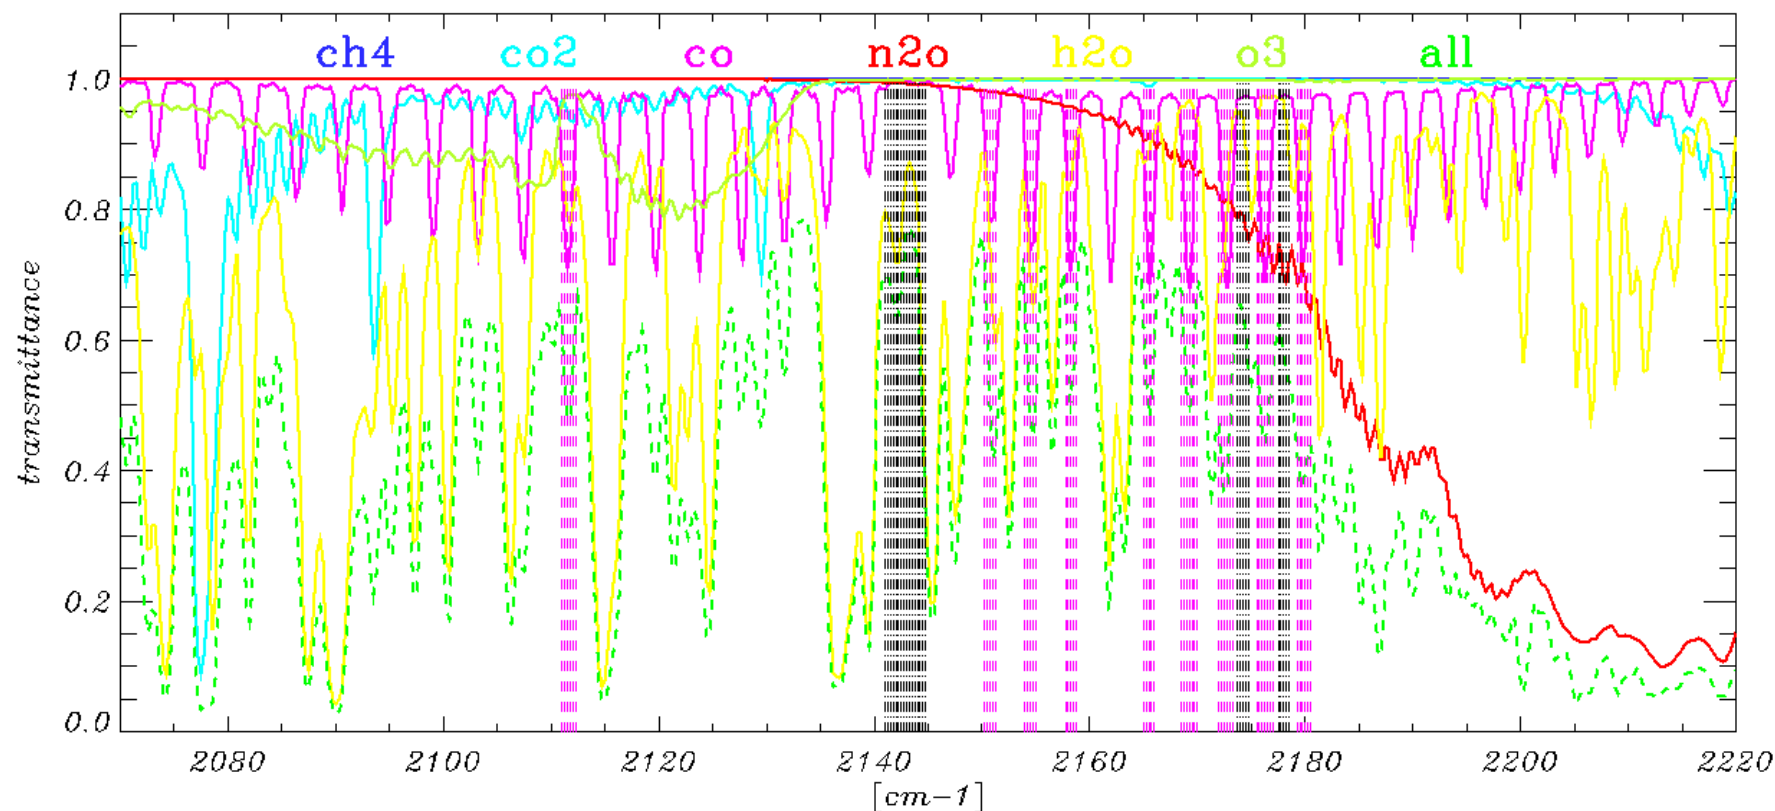


Figure 6: Modified channel selection. Pink and black vertical dashed lines highlight the CO absorption and baseline channels involved in the ANN in the proposed upgrade. Transmittances of various atmospheric components in the CO spectral region were computed for a tropical case.




Figure 7: Synaptic weights between the input (horizontally) and the first hidden layer (vertically) after training for CO total column retrieval (test configuration 0ah). From left to right, the inputs are respectively the absorption lines (labels 0 to 29, blue), the baseline radiances (0 to 2, red), the surface temperature (T_s), a temperature profile (0 to 17, black), the satellite zenith angle (z), the surface pressure (P_s) and the solar zenith angle (S).




Figure 8: *Jacobians of the ANN CO total column retrieval computed over the Indian Ocean on 28 August 2008. The various IFOVs (1000 in total) are stored along the horizontal axis while the derivatives to the inputs are ordered on the vertical axis from bottom to top in the same way as they are listed in Figure 7. The four last inputs (Aux) are respectively the satellite zenith angle (secant and cosine), the surface pressure and the solar zenith angle (cosine).*

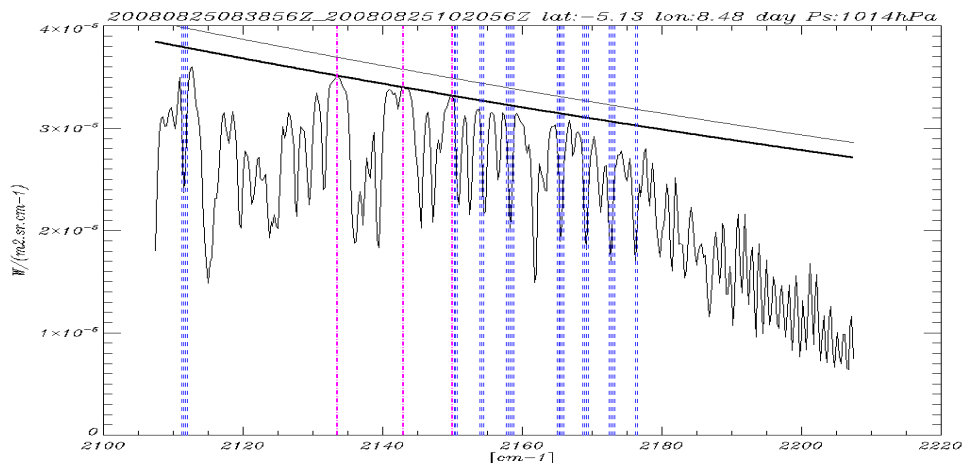


Figure 9: CO spectral region measured off the coasts of Gabon on 25 August 2008. Channels originally selected by SA are shown in blue (absorption lines) and pink (baseline). The thermal background plotted with a thick line is the Planck function computed after the baseline channels while the thinner envelope was computed with the surface temperature and the theoretical surface emissivity of the scene.

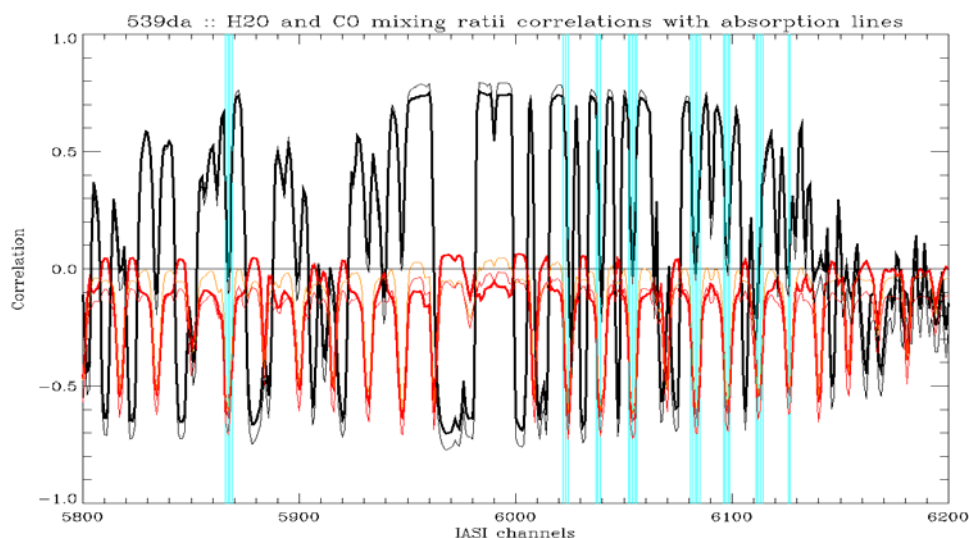


Figure 10: Correlation of line intensities with tropospheric CO (red) and H₂O (black) mixing ratio at various levels in the spectral CO region. The cyan lines indicate the channels selected by SA and involved in the ANN CO retrieval. 53,980 cases from the synthetic data set 539da were used here.

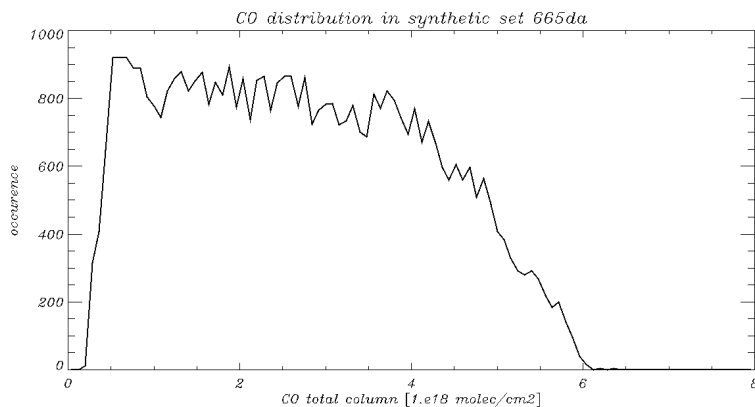


Figure 11: Distribution of the CO total column in the selection of synthetic cases for ANN retrieval training

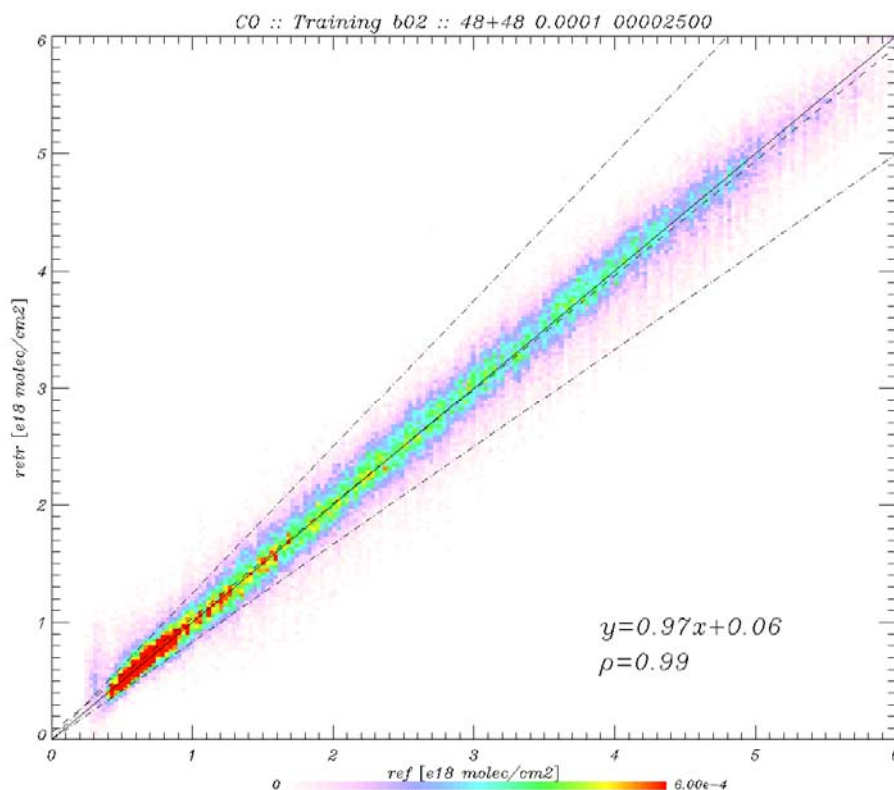


Figure 12: Linear fit on the retrieved CO total columns vs. their associated target after training with noisy inputs. Dash-dot lines denote the $\pm 20\%$ EURD thresholds.

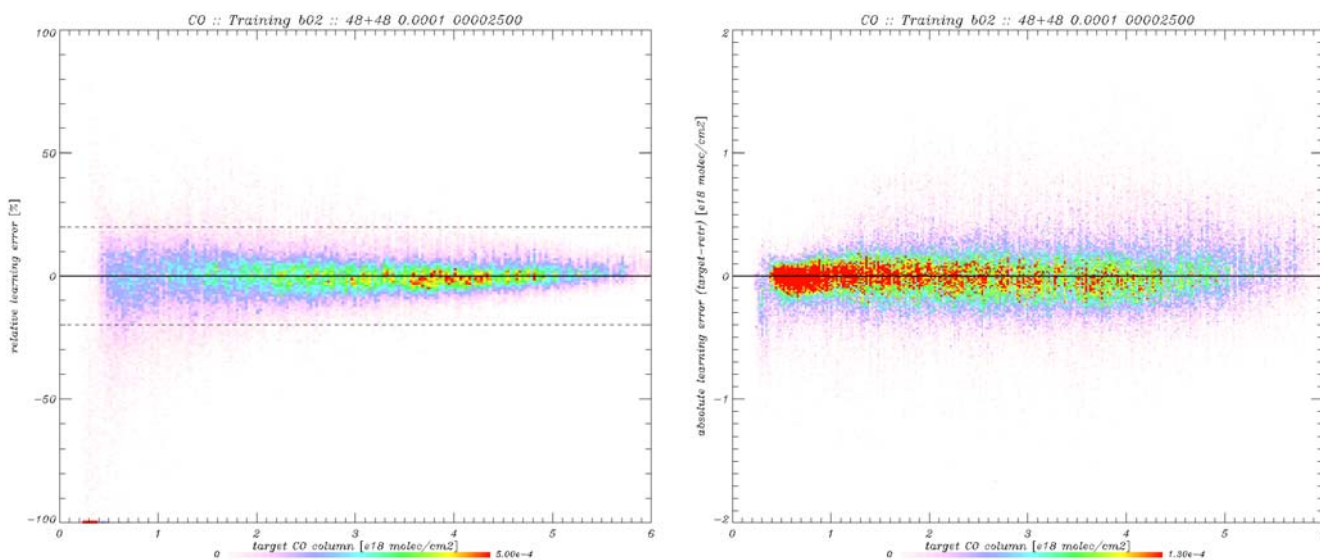


Figure 13: CO training error (left: relative error in %, right: absolute error in 10^{18} molec/cm²) as a function of the target CO column. Noisy inputs were used for training. Dashed lines at $\pm 20\%$ show the thresholds after EURD.

CO :: Training b02 :: 48+48 0.0001 00002500





Figure 14: CO absolute (left) and relative (right) training errors. The overall statistics are displayed in black while the best fitting Gaussian distributions and associated numbers are shown in red. Below, the 7% faintest columns ($CO < 0.7 \text{ molec/cm}^2$) were excluded and statistics recomputed. Artificial networks were taught with noisy inputs.

CO :: Training b02 :: 48+48 0.0001 00002500



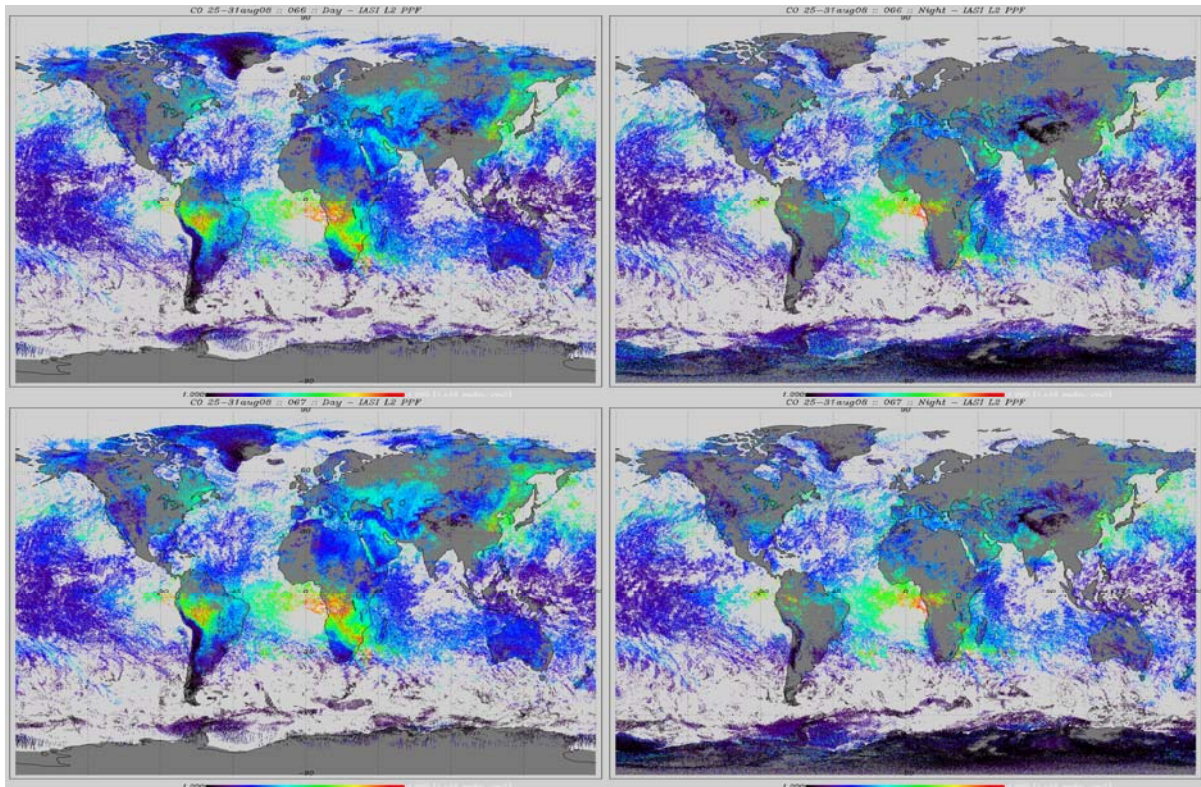


Figure 15: Accumulation of IASI ANN CO retrievals for the 25-31 August 2008 period. Day and night maps are shown on the left- and right-hand side respectively. In configuration 067 (second line), the input temperatures were taken from EOF retrieval (T_s) and ECMWF analyses (vertical profile), whereas in configuration 066 (first line) they all came from EOF retrievals to reproduce what is done in operations.




Figure 16: CO L3 weekly products computed after IASI ANN CO retrievals (top and middle) and MOPITT CO products (bottom) for the period 25-31 August 2008. Day and night maps are shown on the left- and right-hand side respectively. In configuration 067 (middle line), the input temperatures were taken from EOF retrieval (T_s) and ECMWF analyses (vertical profile), whereas in configuration 066 (top line) they all came from EOF retrievals to reproduce what is done in operations.

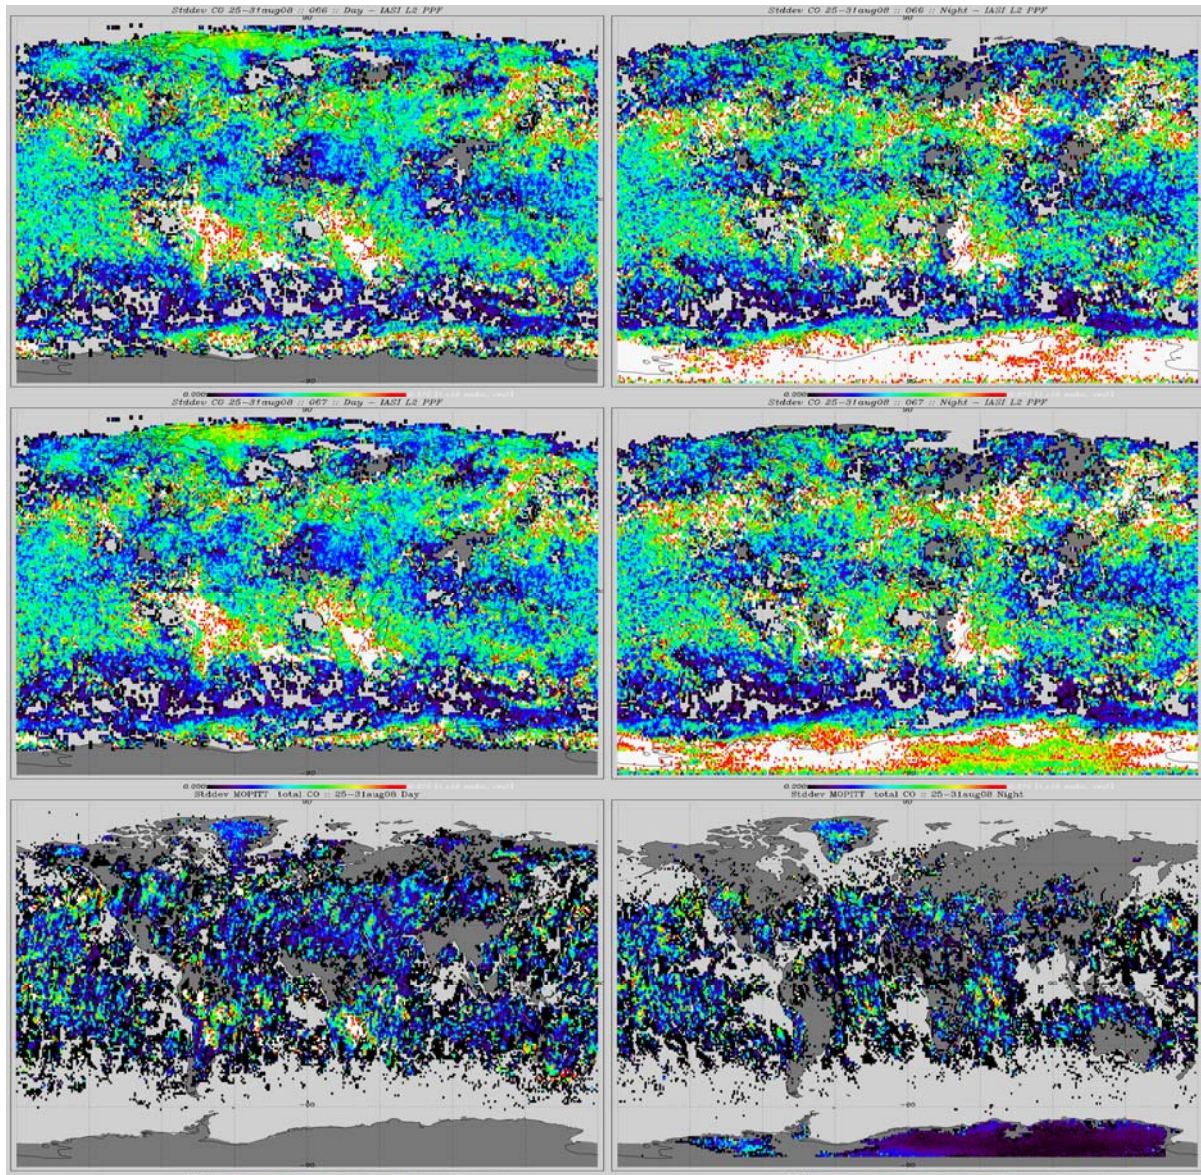


Figure 17: Standard deviation of CO total columns retrieved during the last week of August 2008 with IASI ANN CO (top and middle) and MOPITT CO products (bottom). Day and night maps are shown on the left- and right-hand side respectively. In configuration 067 (middle line), the input temperatures were taken from EOF retrieval (T_s) and ECMWF analyses (vertical profile), whereas in configuration 066 (top line) they all came from EOF retrievals to reproduce what is done in operations.




Figure 18: Angular variation of the mean CO total column computed with IASI measurements between 25 and 31 August 2008 under various ANN configurations. + and • symbols correspond to night and day time, respectively. 013b corresponds to the version in operation since the last PPF upgrade in January 2009. 066 and 067 result from the same network, proposed for the next upgrade, respectively fed with EOF retrieved and ECMWF analyses temperature profiles.




Figure 19: ANN CO retrieval sensitivity to the skin surface temperature: mean CO columns (top) and standard deviations (bottom) were computed for each scan position (EFOV) for a 7-day period in August 2008. + and • symbols correspond to night and day time, respectively. Network 067 (right) differs from 052 (left) by the inclusion during its training of some errors modelling the input temperature characteristics and an upgraded channel selection.

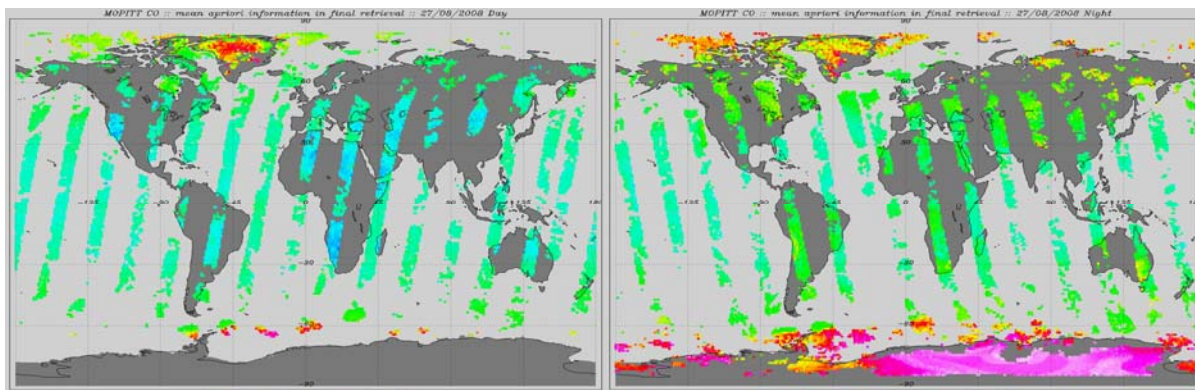


Figure 20: Mean fraction of the a priori information contained in the final MOPITT CO profiles on 27 August 2008 for day (left) and night (right) times




Figure 21: MOPITT CO averaging kernels (<http://mopitt.eos.ucar.edu>)

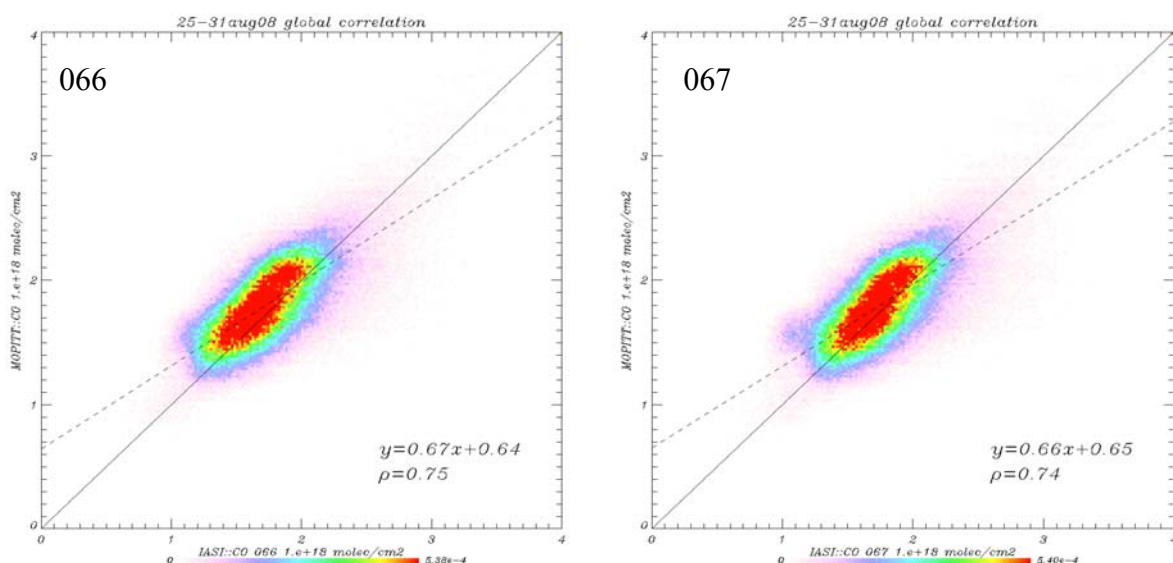


Figure 22: Global correlation between IASI and MOPITT CO retrievals in the last week of August 2008. EOF and ECMWF temperature profiles were used as inputs to the ANN in the configurations 066 and 067, respectively.

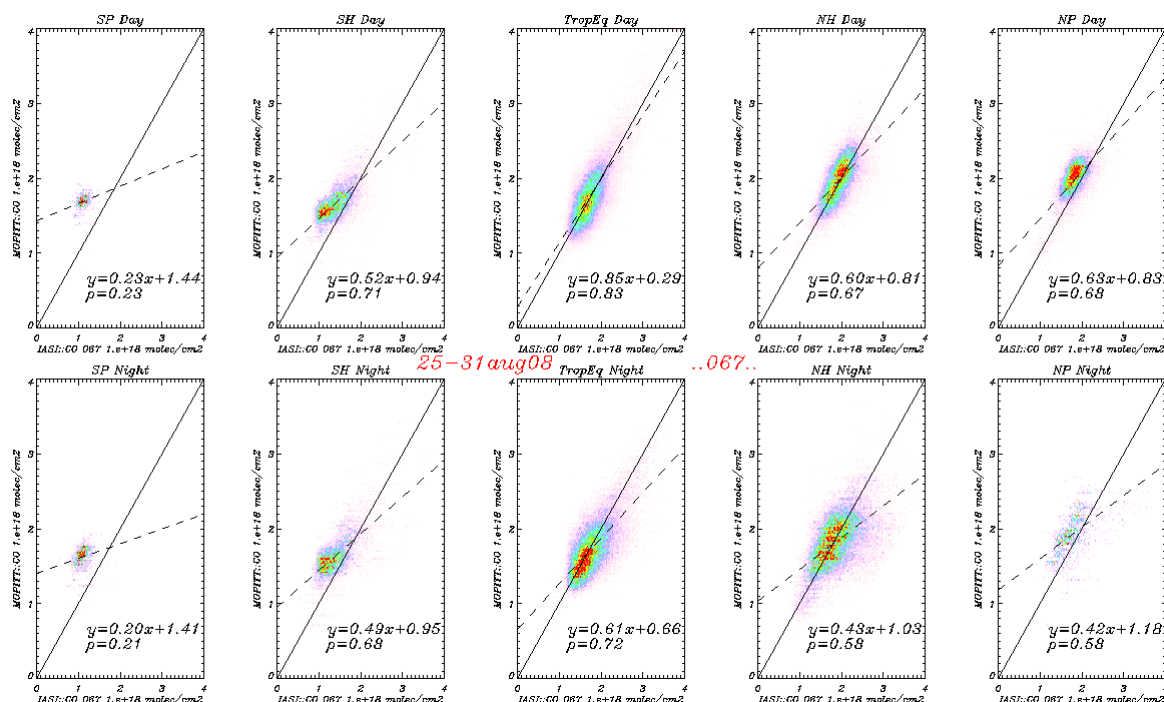
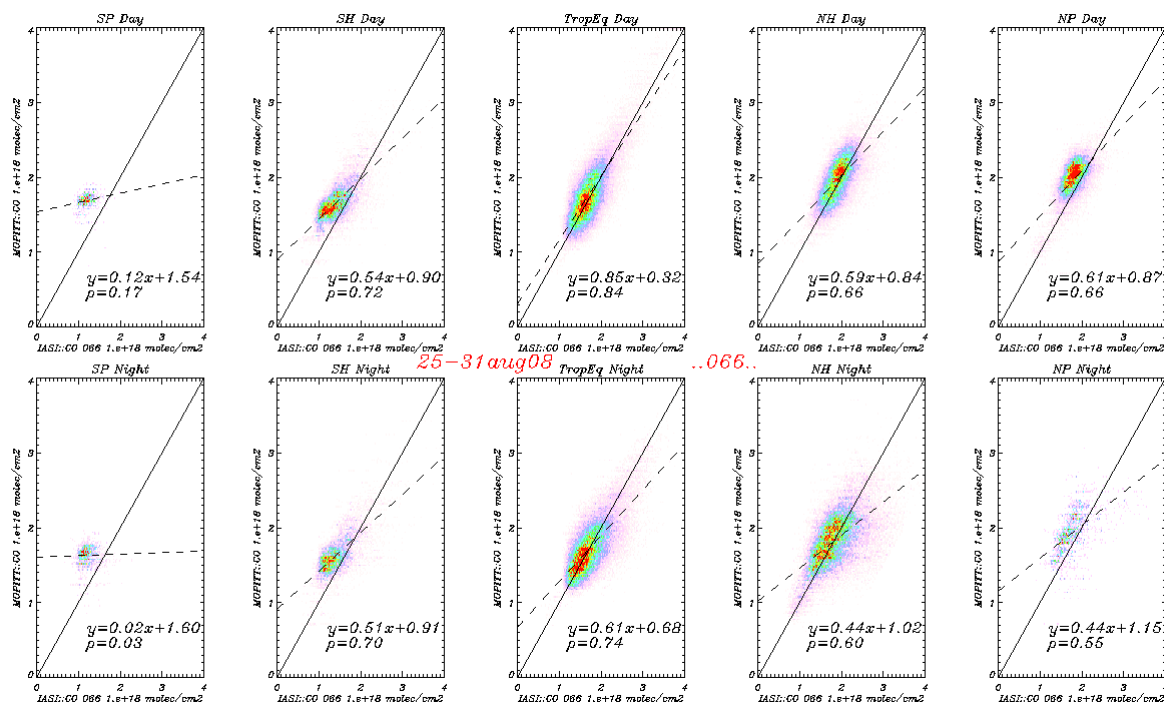


Figure 23: Day/night and geographical breakdown of the correlation between IASI and MOPITT CO averages retrieved for the last week of August 2008. ECMWF and EOF temperature profiles were used as inputs to the ANN in configurations 067 (above) and 066 (below), respectively. From left to right the latitudes are grouped in the following bands: SP [-90° ; -60°], SH [-60° ; -30°], TropEq [-30° ; 30°], NH [30° ; 60°] and NP [60° ; 90°].



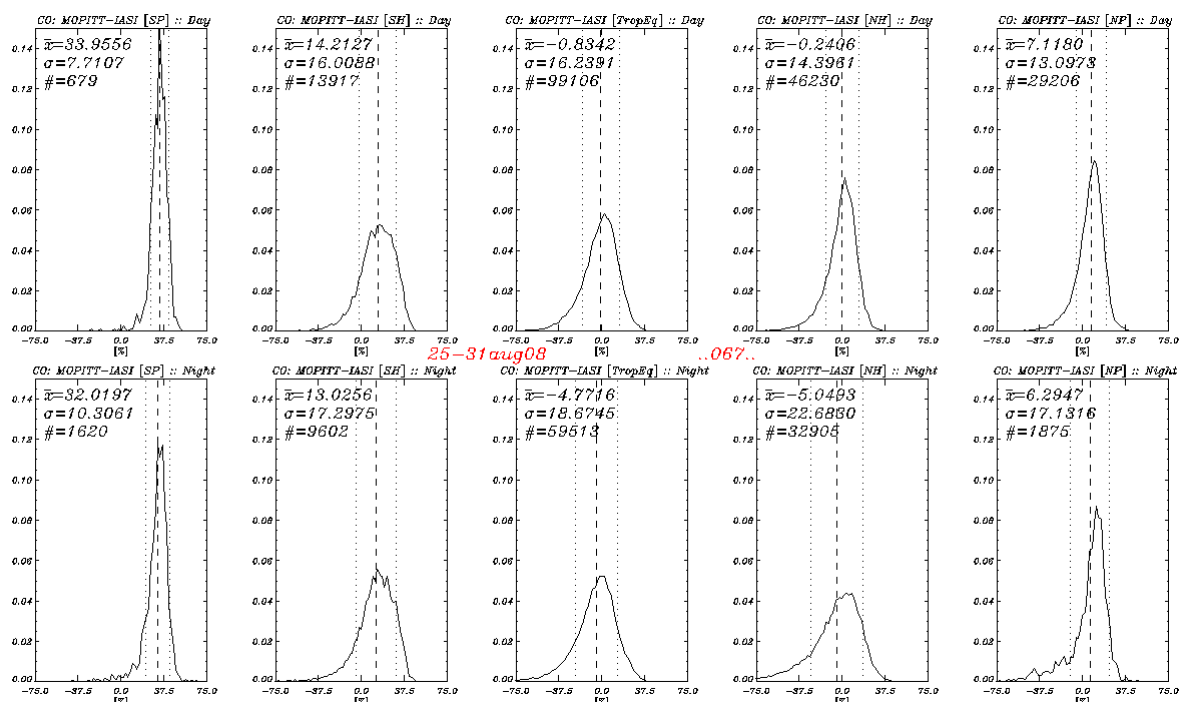



Figure 24: Day/night and geographical breakdown of the IASI relative departure from MOPITT CO averages retrieved for the last week of August 2008. ECMWF and EOF temperature profiles were used as inputs to the ANN in configurations 067 (above) and 066 (below), respectively. From left to right the latitudes are grouped in the following bands: SP [-90° ; -60°], SH [-60° ; -30°], TropEq [-30° ; 30°], NH [30° ; 60°] and NP [60° ; 90°].



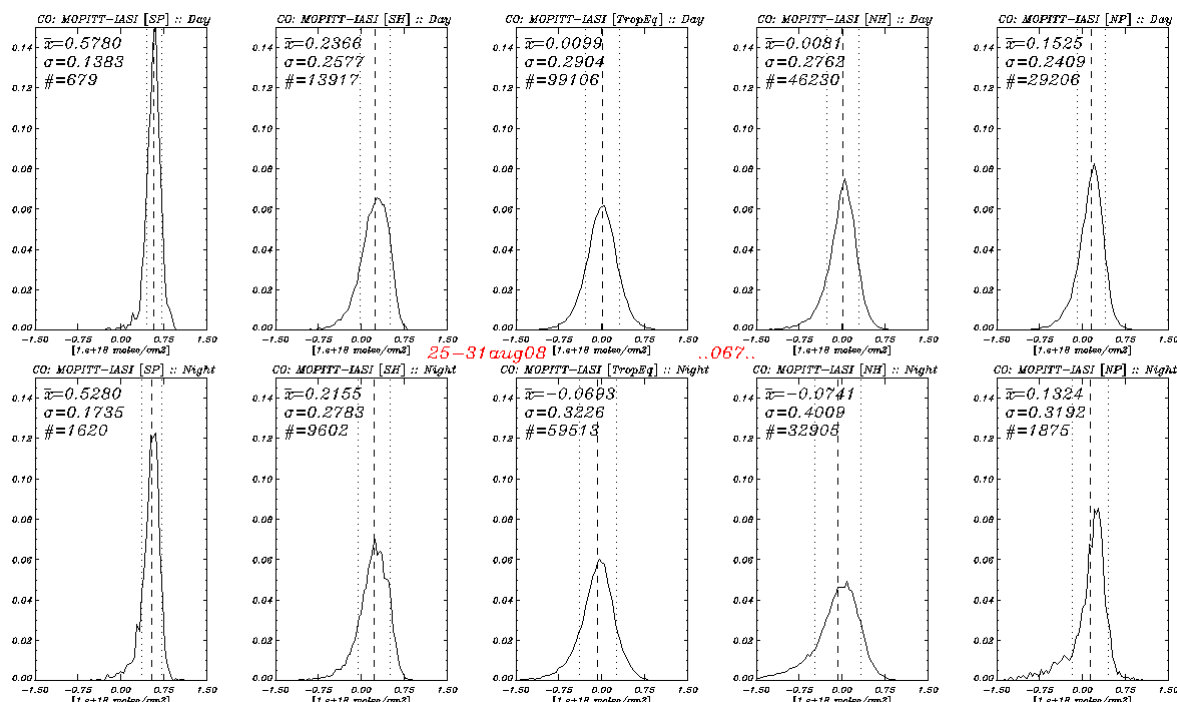
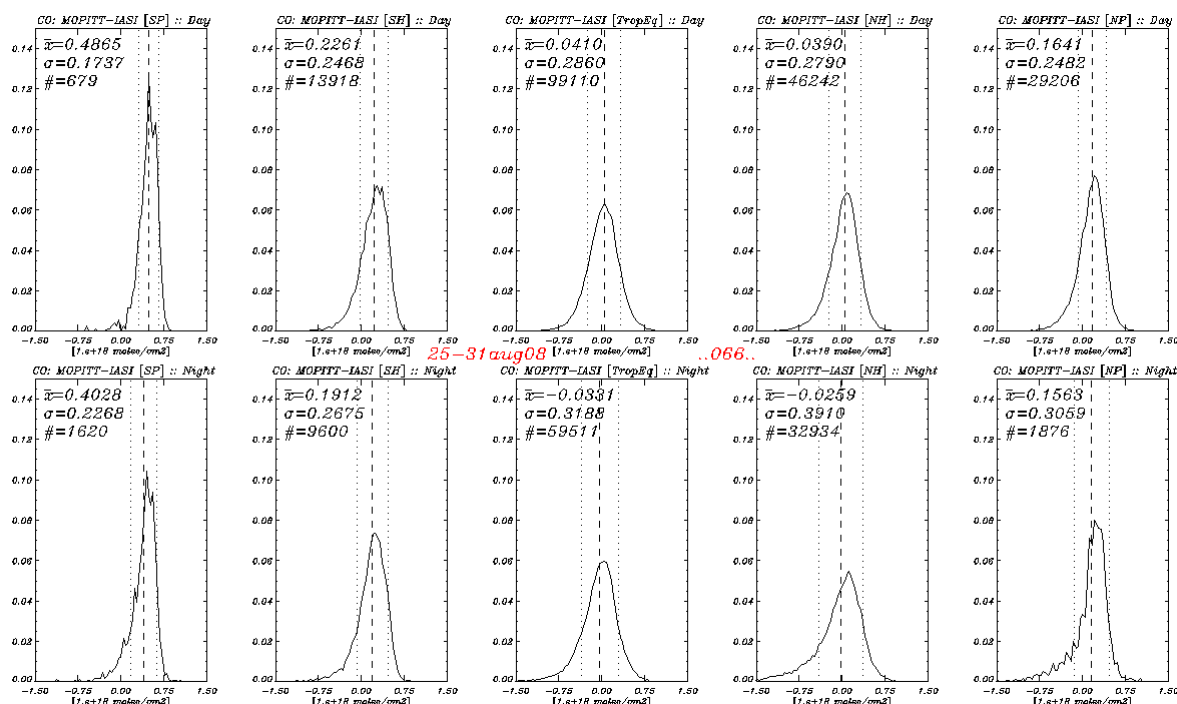


Figure 25: Day/night and geographical breakdown of the IASI absolute departure from MOPITT CO averages retrieved for the last week of August 2008. ECMWF and EOF temperature profiles were used as inputs to the ANN in configurations 067 (above) and 066 (below), respectively. From left to right the latitudes are grouped in the following bands: SP [-90° ; -60°], SH [-60° ; -30°], TropEq [-30° ; 30°], NH [30° ; 60°] and NP [60° ; 90°].






Figure 26: Linear fit on the retrieved CO total columns vs. their associated target after training with noise-free inputs. Dash-dot lines denote the $\pm 20\%$ EURD thresholds.




Figure 27: CO training error (left: relative error in %, right: absolute error in 10^{18} molec/cm²) as a function of the target CO column. Noise-free inputs were used for training. Dashed lines at $\pm 20\%$ show the EURD thresholds.

CO :: Training 0a0 :: 24+24 0.0001 00010000





Figure 28: CO absolute (left) and relative (right) training errors. The overall statistics are displayed in black while the best fitting Gaussian distributions and associated numbers are shown in red. Below, the 7% faintest columns ($CO < 0.7 \text{ molec/cm}^2$) were excluded and statistics recomputed. Artificial networks were taught with noise-free inputs.

CO :: Training 0a0 :: 24+24 0.0001 00010000



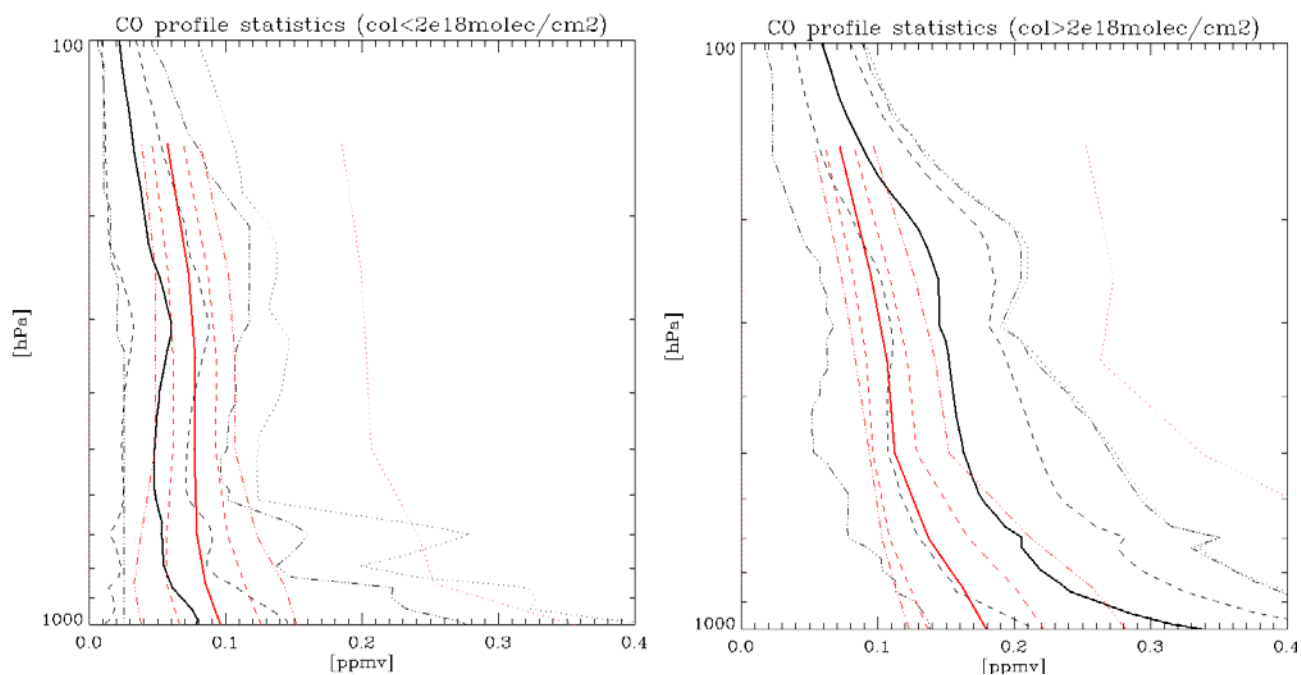


Figure 29: Statistics of CO profiles in the synthetic data sets (black) and in MOPITT CO products for 2008 (red). The central thick plain lines show the mean profile while the dashed, dot-dash and dot lines respectively indicate the ($\text{mean} \pm 1\sigma$), the (2.5%-97.5%) and the (min ; max) limits of these CO vertical distributions.

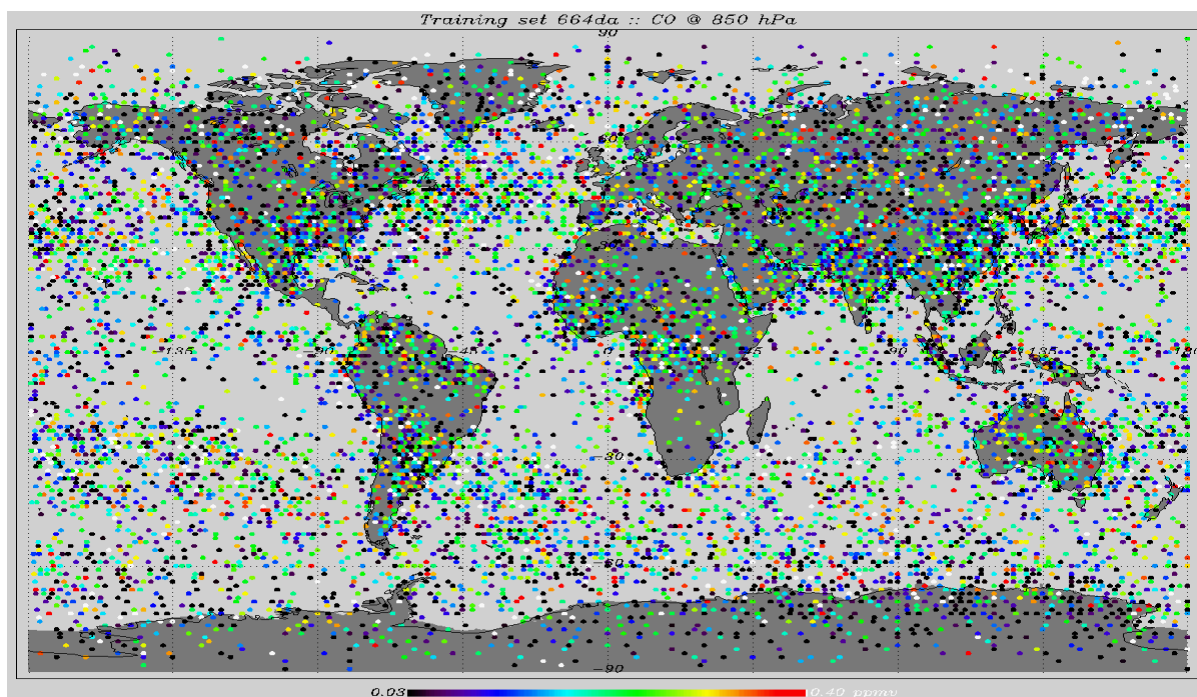


Figure 30: CO mixing ratio around 850 hPa in the synthetic training set showing the geographical distribution of the samples




Figure 31: CO profile correlation as found in MOPITT L3 daily products for the year 2008 (left) and in the synthetic training set (right). The pressure levels are indicated in hPa.




Figure 32: Full CO profile correlation in the synthetic data set represented on the RTIASI grid. The bottom of the atmosphere is at the lower left corner and the TOA at the upper right.

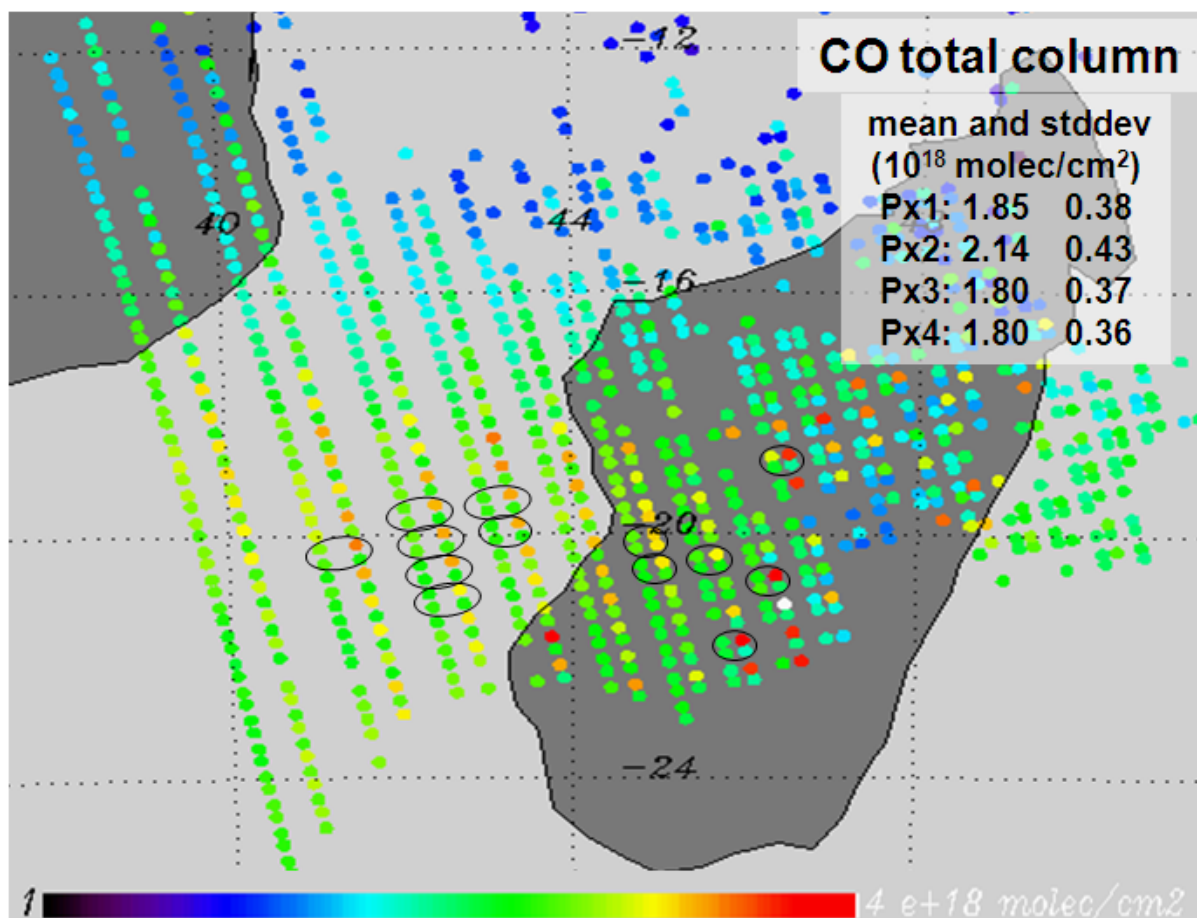


Figure 33: IASI L2 CO products over Madagascar illustrating the outlying pixel 2. Some EFOVs were circled for convenience to identify this pixel (the upper right one). The mean and standard deviation of the CO total column were computed with the entire orbit 20070921171457Z_20070921185657Z (2007 September 21 17:14:57-18:56:57 UTC).




Figure 34: Mean radiance differences between pixel 3 and pixels 1 (blue), 2 (brown) and 4 (green) computed for the scene shown in Figure 33




Figure 35: Radiance biases between pixels (1,3,4) and pixel 2 computed between May 2007 and December 2009 in the IASI L0/L1 NRT monitoring [RD 15] compared to the requirement for interpixel differences (dashed lines)




Figure 36: Same as Figure 17 top line but with pixels 3 and 4 only

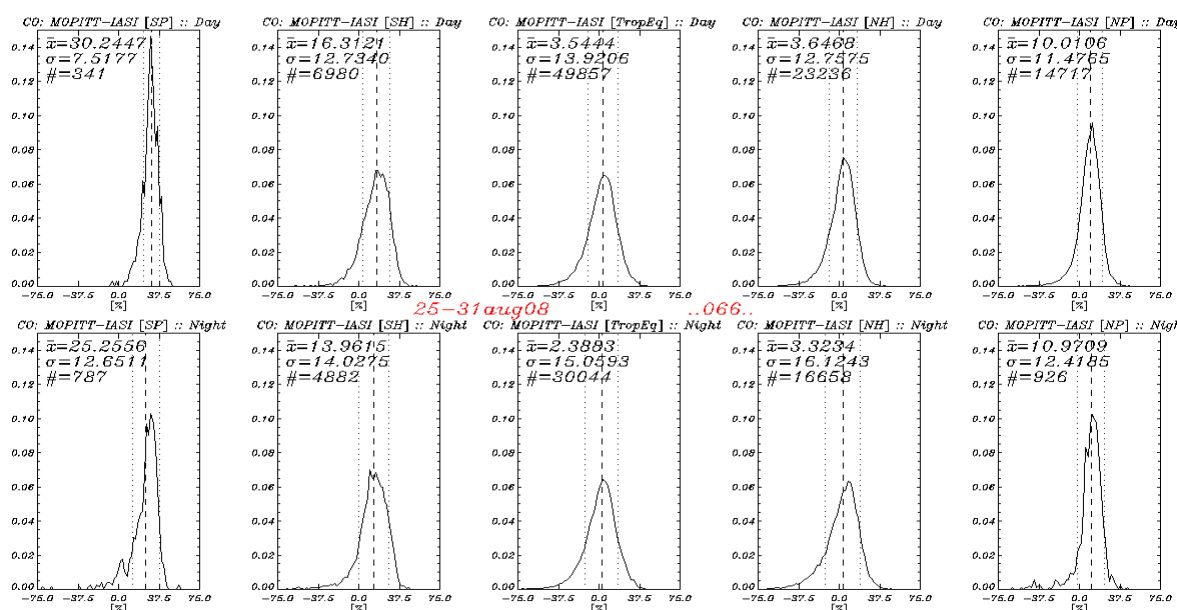


Figure 37: Same as Figure 24, with pixels 3 and 4 only

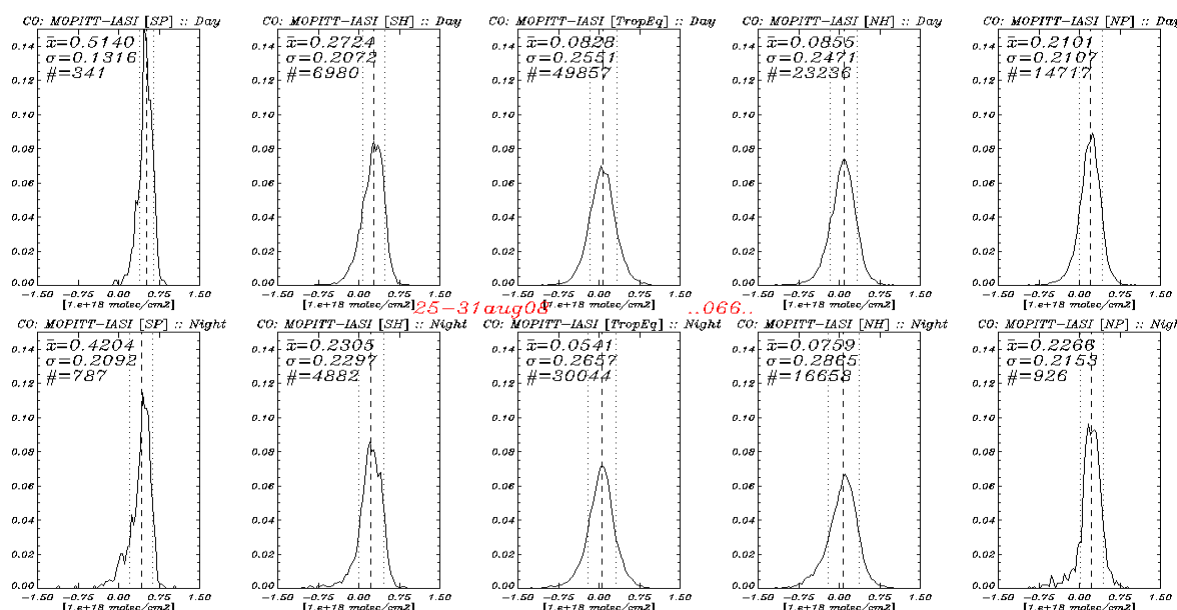


Figure 38: Same as Figure 25, with pixels 3 and 4 only




Figure 39: Monthly mean CO total columns computed with the IFOVs 3 and 4 for the months of August 2008 (top), November 2008 (middle) and February 2009 (bottom) at day (left) and night (right) times

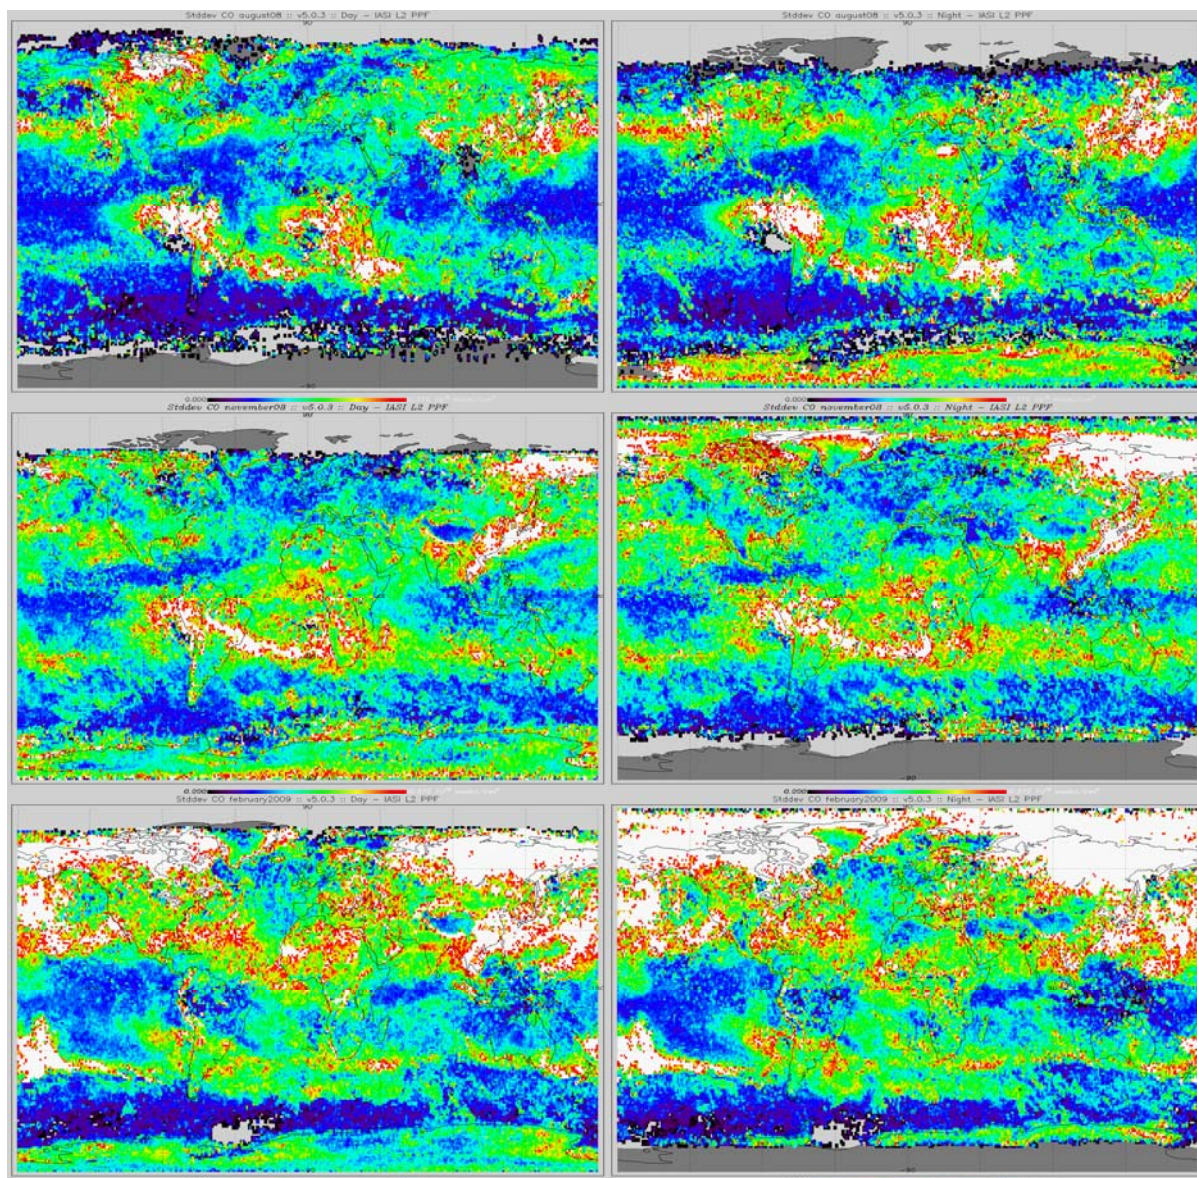


Figure 40: Standard deviations in gridded monthly CO total columns computed with the IFOVs 3 and 4 for the months of August 2008 (top), November 2008 (middle) and February 2009 (bottom) at day (left) and night (right) times





Figure 41: Same as Figure 22-left, with pixels 3 and 4 only and for the months of August 2008 (above) and November 2008 (below)







Figure 42: Same as Figure 23-bottom, with pixels 3 and 4 only and for the months of August 2008 (above) and November 2008 (below)




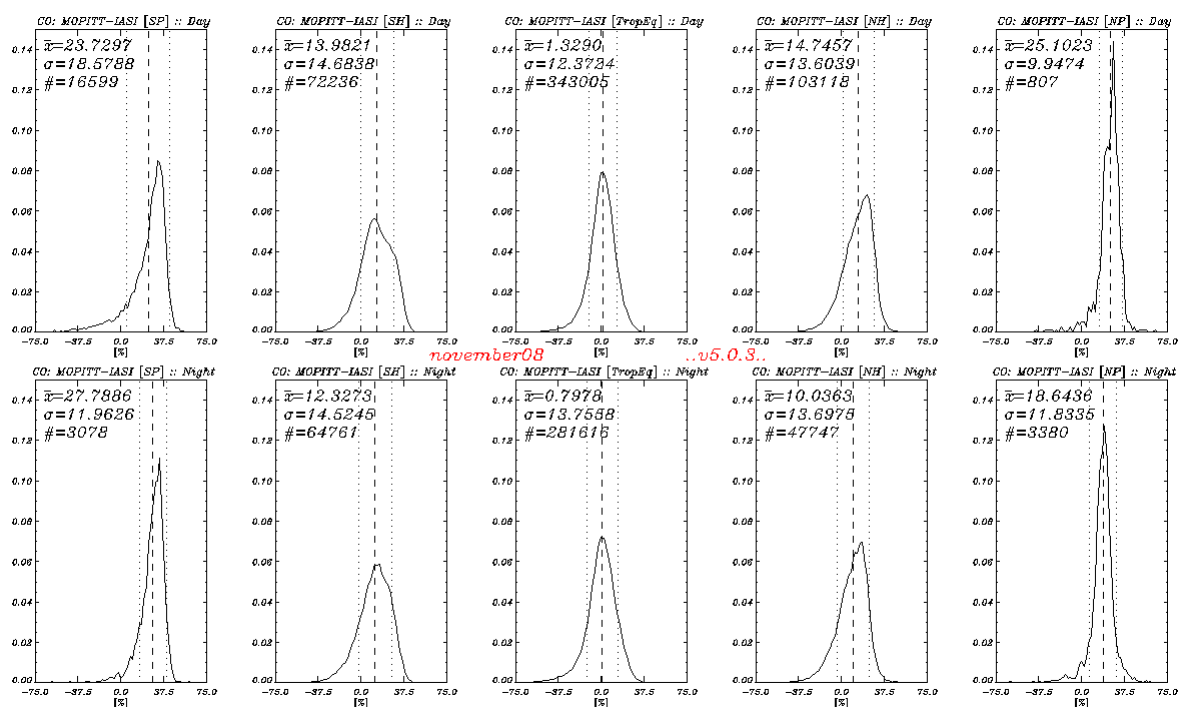


Figure 43: Same as Figure 24-bottom, with pixels 3 and 4 only and for the months of August 2008 (above) and November 2008 (below)






Figure 44: Same as Figure 25-bottom, with pixels 3 and 4 only and for the months of August 2008 (above) and November 2008 (below)

