# ENSAI: AN EMULATOR FOR ATMOSPHERIC CHEMICAL ENSEMBLES

A PREPRINT

Michael Sitwell

Air Quality Research Division Environment and Climate Change Canada Toronto, Ontario, Canada michael.sitwell@ec.gc.ca

April 23, 2025

#### ABSTRACT

Ensemble-based methods for data assimilation and emission inversions are a popular way to encode flow-dependency within the model error covariance. While most ensemble methods do not require the use of an adjoint model, the need to repeatedly run a geophysical model to generate the ensemble can be a significant computational burden. In this paper, we introduce *EnsAI*, a new AI-based ensemble generation system for atmospheric chemical constituents. When trained on an existing ensemble for ammonia generated by the GEM-MACH air quality model, it was shown that the ensembles produced by EnsAI can accurately reproduce the meteorology-dependent features of the original ensemble, while generating the ensemble 3,300 times faster than the original GEM-MACH ensemble. While EnsAI requires an upfront cost for generating an ensemble used for training, as well as the training itself, the long term computational savings can greatly exceed these initial computational costs. When used in an emissions inversion system, EnsAI produced similar inversion results to those in which the original GEM-MACH ensemble was used while using significantly less computational resources.

# 1 Introduction

Ensemble forecasting of atmospheric models has become an invaluable tool in atmospheric modeling over the past few decades. An ensemble of atmospheric states has the ability to represent the uncertainty associated with an atmospheric model and can encode information about both the level of uncertainty of particular state variables as well as the correlation between the uncertainties of different variables. Quantifying the uncertainty of atmospheric forecasts is not only useful when interpreting the results of the forecast, but is also essential for data assimilation and other inversion methods that blend the forecast with observational information.

The goal of an assimilation process is to produce what is know as the *analysis*, which is a combination of forecast (or *background*) and observational information in a way that is optimal (or close to optimal) according to some metric. The uncertainties ascribed to both the model and observations help determine this optimal combination. The error covariances associated with the model and observations set both the level of impact that the observations have on the analysis as well as the spread of the observational information via the error correlations. The error correlations not only spread observational information spatially, but can also impose balance between different variables in multivariate applications (Bannister, 2008). Furthermore, as the analysis increment lies in the subspace spanned by the background error covariance with sufficient rank is essential.

While having an accurate estimate of a model's error covariance is of high importance in many applications, actually obtaining such an estimate is often difficult. Earlier numerical weather prediction (NWP) assimilation systems often used static forecast error covariances as this was more computationally feasible at the time. But as more computational resources became available, more weather centers were able to incorporate flow-dependent error covariances within

their systems. One popular method for doing this was by using four-dimensional variational (4DVar) (Thepaut and Courtier, 1991; Zupanski, 1993; Rabier et al., 2000). However, 4DVar's need for a model adjoint, which can be taxing to develop and maintain, has increased the popularity of adjoint-free methods, in particular algorithms that use ensembles to encode flow-dependent information. The most common ensemble-based algorithms currently employed at weather centers are the ensemble Kalman filter (EnKF) (Evensen, 1994; Lorenc, 2003; Houtekamer and Zhang, 2016) and ensemble–variational assimilation (EnVar) (Buehner et al., 2010; Desroziers et al., 2014; Fairbairn et al., 2014).

While ensemble-based assimilation methods offer many benefits over other methods, a significant downside of using these methods is the computational cost of creating these ensembles, which often requires running an atmospheric model tens to hundreds of times for each time step. The computational cost often limits the number of ensemble members, which can result in spurious correlations within the covariance.

The assimilation of atmospheric chemical observations has been an important component of air quality modeling (Miyazaki et al., 2012b; Robichaud, 2017; Huijnen et al., 2020; Park et al., 2021; Garrigues et al., 2022), stratospheric ozone and UV forecasting (Lahoz et al., 2007; Inness et al., 2013; Rochon et al., 2019), climate modeling (Peters et al., 2010; Miyazaki et al., 2011; Bruhwiler et al., 2014; Agustí-Panareda et al., 2023), and emissions inversions (Miyazaki et al., 2012b,a; Zhu et al., 2013; Cao et al., 2020). The computational challenges associated with ensemble generation are compounded when applied to atmospheric chemistry models, which typically require more computational resources than NWP models. While many operational NWP ensemble system are run at weather centers with ensemble sizes on the order of tens to hundreds, the higher computational requirements of running atmospheric chemistry models make running a similarly designed chemical data assimilation system difficult, especially in an operational setting. When ensemble-based chemical data assimilation is feasible to run, it typically has a much smaller ensemble size than in NWP systems, often with less than a hundred members (Pendergrass et al., 2023). In contrast to most modern NWP systems, it is not uncommon for chemical data assimilation systems to use static background error covariances. For example, at Environment and Climate Change Canada (ECCC), the two operational systems that produce analyses for chemical fields, the air quality surface analysis (Ménard and Deshaies-Jacques, 2018a,b,c) and the stratospheric ozone forecasting system (Rochon et al., 2019; Buehner et al., 2024), both use static background error covariances. ECCC's air quality forecasting system uses the Global Environmental Multiscale - Modelling Air quality and Chemistry (GEM-MACH) chemical weather model, which takes approximately four to five times more computing time and up to nine times more disk space than the Global Environmental Multiscale (GEM) NWP model that GEM-MACH is based on, illustrating the challenges with using ensemble methods with atmospheric chemical models.

The recent explosion in machine learning–based weather prediction (MLWP) (Pathak et al., 2022; Bi et al., 2023; Chen et al., 2023a,c; Lam et al., 2023) offers the potential of atmospheric forecasting that requires a fraction of the run time as compared to traditional physically-based models. The drastic reduction in computational run time is due to both the artificial intelligence (AI) algorithms used as well as running on specialized hardware such as graphics processing units (GPUs). Soon after the initial release of MLWP models, new MLWP systems were developed to produce ensembles of forecasts (Price et al., 2023; Li et al., 2024a), combine data assimilation methods into MLWP models (Xiao et al., 2023; Li et al., 2024b), and end-to-end forecasting/assimilation systems (Chen et al., 2023b; Vaughan et al., 2024). The ability of these MLWP models to meet or exceed the performance (at least with respect to some metrics) of traditional physically-based models that run on central processing unit (CPU) computing clusters, while requiring far less computation time, is opening up a new frontier in the atmospheric sciences.

Currently, most MLWP systems are primarily trained using the ERA5 data set (Hersbach et al., 2020). For atmospheric chemistry applications, additional data sets are needed to provide information on chemical constituents during the training of the AI model. The closest analogue at present to the ERA5 reanalysis for atmospheric chemistry is the Copernicus Atmosphere Monitoring Service (CAMS) (Inness et al., 2019) reanalysis, which only dates back to 2003. Recently, Bodnar et al. (2024) developed an AI geophysical system that is pretrained on NWP data and fine-tuned for atmospheric chemistry using the CAMS reanalysis as well as the CAMS analysis for more recent years. However, at present, the atmospheric chemistry data that could potentially be used for training is much more limited as compared to NWP, particularly for ensemble information for chemicals in the troposphere important for air quality forecasting.

The aim of this work is to develop an AI-based model for efficient ensemble generation of tropospheric chemical constituents for the purpose of air quality forecasting, which we will refer to as *EnsAI*. As finding an appropriate training data set for air quality applications is more difficult as compared to NWP, especially for ensemble information, leveraging any existing ensemble air quality data sets for training may be beneficial for the development of AI-based air quality systems. Sitwell et al. (2022) detailed the development of an ensemble-based ammonia emissions inversion system that was used with the GEM-MACH model. As ensembles of atmospheric chemistry models are computationally expensive to generate and this ensemble for ammonia had already been generated, this ensemble data set will be used as a training set for our AI-based ensemble generation system. While generating ensembles for ammonia

emissions inversions will be used as a test application for the AI-generated ensembles, another goal of this work is to develop an AI system flexible enough to be used in the future for generating ensembles for other chemical constituents relevant for air quality.

For our application, the ensemble of emission perturbations is computationally inexpensive to create, but translating the ensemble of emissions perturbations into an ensemble of concentration perturbations using a physically-based atmospheric model is usually very computationally demanding. EnsAI uses a U-Net convolutional neural network to map ammonia emission perturbations to ammonia surface concentration perturbations, using surface wind and temperature fields as auxiliary inputs to the neural network. EnsAI can generate the ensemble of concentrations without needing to run computationally expensive physically-based atmospheric models. As EnsAI has the upfront cost of generating an ensemble for training, as well as the training itself, the intent is to be able to use EnsAI for much longer time periods than the time period used for training, so that these initial costs are much smaller than the long term computational savings over the intended use period.

# 2 Emissions Inversions

This study uses the application to ammonia emission inversions as a testbed for developing the foundations of an AIbased ensemble generation system. As in Sitwell et al. (2022), ammonia retrievals from May to August 2016 from the satellite-borne CrIS instrument are used in emission inversions to refine the ammonia emissions used with the GEM-MACH air quality model. While the inversion methods are described in full in Sitwell et al. (2022), we summarize the main features in this section, as well as a few modifications made for this study.

### 2.1 The GEM-MACH Air Quality Model

GEM-MACH (Moran et al., 2010; Gong et al., 2015; Pavlovic et al., 2016) is a chemical transport model used in ECCC's operational air-quality forecasting system that is embedded within ECCC's operational GEM weather forecasting model (Côté et al., 1998b,a; Girard et al., 2014). GEM-MACH is a physically-based model of atmospheric gas-phase, aqueous-phase, and heterogeneous chemistry. Version 3.1.0 of GEM-MACH was used for this study, which has 85 vertical levels that extend from the surface to 0.1 hPa. GEM-MACH can be run with both global and limited area horizontal grids. For this work, a limited area horizontal grid with spacing of  $0.09^{\circ}$  (~10 km) was used.

This GEM-MACH version also includes an updated version of a bidirectional flux model for ammonia (Whaley et al., 2018) that combines time-independent emission potentials with a resistance-based deposition model to yield meteorologically-dependent ammonia emissions. The meteorology is refreshed at the beginning of each forecast using analyses from ECCC's Global Deterministic Prediction System (Buehner et al., 2015).

Monthly-mean emissions, originating from either an emissions inventory or an emissions inversion, are translated into hourly emissions that can be used in GEM-MACH through the use of sector-dependent time profiles. The unperturbed set of emissions was version 3.1.2 of the operational GEM-MACH emissions data set. Figure 1 shows this emissions inventory for ammonia averaged over May to August, which are the months of interest for this study.



Figure 1: Unperturbed ammonia emissions inventory averaged over May to August. Emissions are only displayed within the domain of the EnsAI model. Particular locations examined in this work are also displayed, which are located in (in order of descending emissions) (a) North Carolina, (b) Iowa, (c) Ohio, (d) Pennsylvania, and (e) Georgia.

#### 2.2 CrIS Ammonia Retrievals

The CrIS instrument, which is aboard the Suomi National Polar-orbiting Partnership (SNPP), NOAA-20 and NOAA-21 satellites, is a Fourier transform spectrometer that measures the infrared spectrum in three different spectral bands. The CrIS Fast Physical Retrieval (CFPR) algorithm (Shephard and Cady-Pereira, 2015; Shephard et al., 2020), based on the TES ammonia retrieval algorithm (Shephard et al., 2011), used for the ammonia retrievals makes use of the longwave band between 960 and 967 cm<sup>-1</sup>. As in Sitwell et al. (2022), we will focus on the spring and summer of 2016 (which comprises most of the peak ammonia emissions season) and use retrievals from the CrIS instrument aboard the SNPP satellite. The SNPP satellite, which is in a sun-synchronous orbit, has local overpass times at approximately 01:30 and 13:30, though only the day-time retrieval is used in the inversions. CrIS has a resolution at nadir of around 14 km. The retrievals are made on 14 pressure levels, with the averaging kernel typically peaking at or near the surface.

Retrievals were from version 1.6.4 of the CFPR algorithm that, in contrast to version 1.5 that was used in Sitwell et al. (2022), includes non-detects when the signal is below the detection limit (White et al., 2023). Only retrievals over land with a signal-to-noise ratio of at least one and have at least 0.1 degrees of freedom are used in the inversions.

The goal of this work is to develop an ensemble generation AI model that runs on GPUs, which typically have access to less memory than traditional CPU clusters. As such, to reduce memory requirements, only the retrieval at the surface will be used in the inversions, instead of all 14 levels that were used in the inversions in Sitwell et al. (2022). This should be a reasonable simplification since the retrievals are mainly sensitive to the surface and lower atmospheric levels.

#### 2.3 Inversion Methods

Inversion algorithms are used to create a statistical blend of observational and a priori (background) information by solving an inverse problem, the result of which is known as the analysis. Both the background model state and observations are weighted by their error covariances. The variances of the error covariances determines the level of impact the observations have on the analysis, while the correlations can (among other things) spread the observational information to other variables in the model state. Most algorithms solve this problem by finding an additive state vector, known as the *increment*, to add to the a priori state vector.

For our current application, a set of emissions  $\mathbf{e}$  is used as input into GEM-MACH to generate atmospheric concentrations  $\mathbf{c}$ , where only  $\mathbf{c}$  is observed. As mentioned in Section 2.1, because we will only be utilizing the lowest level of the model output,  $\mathbf{c}$  will denote only the surface concentrations. If our model state vector  $\mathbf{x}$  is represented in block

form as

$$\mathbf{x} = \begin{bmatrix} \mathbf{c} \\ \mathbf{e} \end{bmatrix},\tag{1}$$

then the background error covariance **B** can be expressed as

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_{cc} & \mathbf{B}_{ec}^{\mathrm{T}} \\ \mathbf{B}_{ec} & \mathbf{B}_{ee}^{\mathrm{T}} \end{bmatrix},\tag{2}$$

where  $\mathbf{B}_{cc}$  and  $\mathbf{B}_{ec}$  represent the univariate covariances for the surface concentrations and emissions, respectively, while  $\mathbf{B}_{ce}$  and  $\mathbf{B}_{ec}$  represent their cross-covariance. Although the uncertainties in the atmospheric concentration of ammonia is likely due to a number of different factors including uncertainties in emissions, meteorology, initial conditions, and chemistry, for these inversions only the uncertainty due to the emissions is considered.

Ensemble methods model the background error covariance, in full or in part, by use of an ensemble of model states, where each member of the ensemble has a different perturbed state. An ensemble  $\{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$  of N model states can be used to define a background error covariance  $\mathbf{B}^{\text{ens}}$  as

$$\mathbf{B}^{\text{ens}} = \frac{1}{N-1} \mathcal{L} \circ \left( \sum_{i=1}^{N} \delta \mathbf{x}_{i} \delta \mathbf{x}_{i}^{\text{T}} \right), \tag{3}$$

where  $\delta \mathbf{x}_i$  is model state perturbation of ensemble member *i*. Physically-based atmospheric models are usually computationally costly to run, so ensembles typically consist of tens to hundreds of members. These small ensemble sizes produce error covariances with low rank and spurious long-distance correlations. A localization scheme is usually included to mitigate these effects, which we implement through an element-wise multiplication (denoted by  $\circ$ ) of the ensemble error covariance with a localization matrix  $\mathcal{L}$ .

As individual CrIS ammonia retrievals have relatively large uncertainties (Shephard and Cady-Pereira, 2015), emissions inversions are generally performed using weeks of retrievals, so that a sufficient amount of observational information is used in each inversion. Consequently, each inversion produces a time-mean ammonia emissions field for the inversion time window. While the inversion produces a single time-mean emissions field, the resulting updated hourly emissions do vary within the inversion time window as the ammonia emission are dependent on meteorological conditions, most notably the surface temperature. Further details of the inversion algorithm can be found in Section A of the Appendix.

#### 2.4 GEM-MACH Ensemble Generation

Ensemble generation is often the most computationally demanding part of an ensemble forecasting system. The training data used in our AI ensemble system was trained on a previously existing ammonia ensemble, which removed the need to generate new training data using a physically-based system. This section briefly summarizes the processes that were used to generate this ensemble.

The ensemble generation process consists of two steps: First, for each ensemble member *i*, a Gaussian random field is drawn that is used to construct the time-mean emissions  $\mathbf{e}_i$ . These emissions are then combined with the unperturbed model input, collective denoted by  $\boldsymbol{\theta}$ , which includes the emissions from all chemical species other than ammonia as well as initial conditions for both meteorological and chemical concentration fields. The GEM-MACH model, here represented by the operator  $\mathcal{M}^{GM}$ , is then run to produce the atmospheric concentrations  $\mathbf{c}_i$  for that ensemble member, which is represented symbolically as

$$\mathbf{c}_i = \mathcal{M}^{\mathrm{GM}}(\mathbf{e}_i, \boldsymbol{\theta}). \tag{4}$$

The initial step of generating the emissions ensemble is relatively computationally inexpensive. A random field is repeated drawn from an isotropic normal distribution and are used to perturb the ammonia emissions. The standard deviations of the emissions perturbation field was set to 50% of the monthly mean values, with a minimum standard deviation value of  $10 \text{ mg s}^{-1} \text{ km}^{-2}$  imposed so that all locations have a non-negligible ensemble variance. The horizontal spatial correlations of this distribution were chosen to have a half-width at half-maximum of 40 km. The second step of repeatedly running GEM-MACH with the perturbed emissions fields accounts for most of the computational cost of this inversion method and is the motivation for exploring AI-based ensemble generation methods that can potentially produce an ensemble at much lower computational cost.

## **3** AI Ensemble Generation

As repeatedly running GEM-MACH is very computationally demanding for ensembles of any reasonable size, we wish to find an AI-base model  $\mathcal{M}^{\mathrm{EnsAI}}$  that can can be used in Eq. (4) in place of  $\mathcal{M}^{\mathrm{GM}}$  that emulates the GEM-MACH model but is able to run much faster than GEM-MACH.

In the original ensemble generation procedure described in Section 2.4, we note from Eq. (4) that the full GEM-MACH model  $\mathcal{M}^{GM}$  is used to produce the *perturbed* concentration fields  $\mathbf{c}_i$ , which in our case are the surface ammonia concentrations. A full GEM-MACH emulator would aim to reproduce the full forecasts for multiple chemical species. However, as our goal is only to reproduce ensemble perturbations,  $\mathcal{M}^{EnsAI}$  will only predict the ensemble perturbations of the ammonia surface concentrations  $\delta \mathbf{c}_i$  given the perturbations to the ammonia emissions  $\delta \mathbf{e}_i$ . Since  $\mathcal{M}^{EnsAI}$  only needs to simulate the ensemble perturbations of certain model fields (only surface ammonia in our case),  $\mathcal{M}^{EnsAI}$  can potentially be much simpler than an emulator that aims to reproduce the full GEM-MACH output.

For application to ammonia emission inversions, the same perturbation to the time-mean ammonia emissions can results in different perturbations to the atmospheric ammonia concentration at different times due to the dependency on the meteorological conditions. The strength of the ammonia emissions are strongly dependent on temperature, while the local spatial distribution of its atmospheric concentration is dependent on advection patterns. This is illustrated in Figure 2, which displays the ammonia emission and surface concentration perturbations for one particular member of the ensemble. The monthly-mean ammonia emission perturbations for May is shown in the top row, while the the bottom row shows the surface concentration perturbations resulting from this emission perturbation at three different times (all at the same time, 18 UTC, but for different days). While all three surface concentration perturbations originate from the same emissions perturbation, the different meteorological conditions at each time have a clear effect on both the magnitude and spatial pattern of surface concentration perturbations.



Figure 2: Monthly-mean ammonia emission perturbations for May (top row) and the resulting ammonia surface concentration perturbations at three different time in May (bottom row) for one member of the GEM-MACH ensemble.

For the ensemble generation described in Eq. (4), the input to GEM-MACH consists of a randomly perturbed ammonia emissions field  $\mathbf{e}_i$ , as well as a large number of unperturbed fields  $\theta$ . In addition to meteorological fields, GEM-MACH requires hundreds of chemical constituent fields as well as other auxiliary fields such as those describing land-use and soil hydrology. As GEM-MACH is typically run on large CPU clusters that have access to much larger amounts of memory than is currently available to GPUs, EnsAI was designed to use drastically fewer unperturbed input fields. As ammonia emissions have a strong temperature dependence and the local spatial patterns of ammonia concentrations

depend on its advection, the unperturbed input fields to EnsAI were limited to the surface temperature **T** and horizontal wind fields **u** and **v**. The EnsAI model can now be summarize as

$$\delta \mathbf{c}_i = \mathcal{M}^{\text{EnsAI}}(\delta \mathbf{e}_i, \mathbf{u}, \mathbf{v}, \mathbf{T}).$$
(5)

Additionally, the original  $772 \times 642$  horizontal grid of the GEM-MACH regional model, which covers all of Canada, the United States, and northern Mexico, was reduced to a  $256 \times 256$  grid that cover only the central and eastern United States and Canada (displayed in Fig. 1) to reduce memory requirements.

It is worth noting that generating the ensemble of emissions is simple, as they are constructed from a Gaussian random field as described in Section 2.4. In Eq. (4), the GEM-MACH model  $\mathcal{M}^{GM}$  translates the emission ensemble into a atmospheric concentration ensemble, and the desire is for  $\mathcal{M}^{EnsAI}$  to emulate this process. This is in contrast to other AI-based geophysical systems, such as those using diffusion-based generative models (Price et al., 2023; Li et al., 2024a), that construct a latent space that is easy to sample in. In diffusion-based generative models, a diffusion process is used to add noise to an image (which in our application would correspond to an atmospheric field) until the image is dominated by noise. The goal of the noising process is for the distribution of noised images to be well described by a probability distribution that is easy to sample from. In these systems, random samples are drawn in the latent space and are then denoised using the reverse diffusion process (Song et al., 2020; Song and Ermon, 2020; Ho et al., 2020) to transform the random samples into the space of the atmospheric model. In our application, sampling from the emissions distribution is already relatively easy, so constructing and transforming into a latent space is not necessary. In both cases, samples are randomly drawn in a space where the probability distribution is relatively simple and the AI model transforms that sample to the state space of the atmospheric model. However, EnsAI can potentially be simpler than diffusion-based (or other similar) generative models since a diffusion process to denoise an image does not need to be modeled.

# **4** Neural Network Architecture

EnsAI is based on the popular U-Net convolutional neural network architecture (Ronneberger et al., 2015) that was originally developed for image segmentation, but has since been widely used for other applications such as diffusionbased image generation (Ho et al., 2020). The U-Net used for this study is displayed in Figure 3, where each box represents either the model state at the input or output, which in the language of image processing can be referred to simply as an *image*, or the image transformed into an intermediate state often referred to as a *feature map*. The U-Net broadly consists of a series of convolutions and max pooling that encodes images into various feature maps, followed by a series of convolutions and upsampling to decode the feature maps back to an image. The dimensions of the horizontal grid at each stage within the U-Net is given on the left of the diagram, while the number of fields, usually referred to as *channels* in the current context, is displayed above each box.



Figure 3: U-Net architecture used for EnsAI. Each box represents an image or feature map, where the x, y dimensions are displayed on the left and the number of channels are displayed above each box. White boxes denote feature maps that have been copied during the encoding sequence and concatenated with feature maps produced in the decoding sequence. The blue, orange, green, and purple arrows represent residual blocks, max pooling, upscaling blocks, and two-dimensional convolution, respectively. This figure is adapted from the depiction of the original U-Net architecture shown in Fig. 1 of Ronneberger et al. (2015).

The input fields to EnsAI are displayed in the upper-left corner of Fig. 3, which consists of the four input fields  $[\delta e, u, v, T]$  on the 256 × 256 horizontal grid. In the encoding branch of the network, the input image is sequentially transformed into feature maps with more channels and a reduction in spatial resolution to encode different image features. The decoding branch decreases the number of channels and increases the spatial resolution, with the final output having the original spatial resolution and a single channel for the surface atmospheric concentration perturbation  $\delta c$  (shown in the upper-right corner of Fig. 3). Details on each component of the U-Net can be found in Appendix B.

## **5** Training the Network

As the existing GEM-MACH ammonia ensemble is used for both training and testing, it must be divided into training, validation, and testing data sets. When multiple years of data are available, it is common to withhold the last year of data as a test set. However, in our case only a single spring/summer period is available that spans from May to August 2016.

Following ECCC's operational air quality system (Ménard and Deshaies-Jacques, 2018c; Ménard et al., 2020), which evaluate model error covariances within two-month time periods, we similarly segment our ensemble into two-month durations: One period for May/June and another for July/August. We withhold the first week of data from June and August, which fall approximately at the mid-point of each two-month time period, to use as our test set. The selected test periods fall approximately after the first and third quarters of the full data set, which should be somewhat more representative of the whole period as compared to selecting the time period at the end of the data set that may be less representative of the spring/summer seasons.

Emissions inversions for ammonia are typically designed so that one inversion is done for each month. Choosing our test set to span only single weeks means that we will only be able to perform inversions using a single week's worth of data as part of our evaluation. However, since our goal is for the EnsAI ensemble to follow the GEM-MACH ensemble as closely as possible, comparing the inversion results from the two different ensembles should be sufficient for our purpose.

A 60 member ensemble of hourly GEM-MACH model output for May to August 2016 was used for training, which amounted to 123 days  $\times$  24 h day<sup>-1</sup>  $\times$  60 members = 177, 120 samples over the whole data set. After removing the test set (about 11% of the total data), an additional 11 days (or about 9% of the total data) was reserved for the validation set used during training. The mean squared error (MSE) was used as the loss function for training. While the MSE loss function can lead to smoothing at smaller scales, no significant smoothing was observed in our

experiments. Training was done on a NVIDIA A100 GPU and took approximately three days to train the network through 50 epochs. More details of the training can be found in Section S1 of the Supplement.

# 6 Run Time Comparison

One of the main motivations of this work is to develop a method for ensemble generation that can run significantly faster than traditional physically-based systems. The original GEM-MACH ensemble took roughly 6.5 hours of wall time to generate a week's worth of simulated data per ensemble member when run on a CPU cluster using 720 computing nodes, with each node having access to 80 cores. In contrast, EnsAI when run on a single GPU took only 7 seconds of wall time to generate a week's worth of simulated data per ensemble member, or about 3,300 times faster than GEM-MACH.

The run time comparisons above for EnsAI do not include the upfront cost of generating the GEM-MACH ensemble used for training as well as the time for training. As previously mentioned, the intent is to run EnsAI for time periods much longer than the training period, so that the long term computational savings of running EnsAI instead of GEM-MACH for ensemble generation becomes much larger than these initial costs. For operational implementation, the desire is to only retrain EnsAI occasionally when there are significant changes in the operational version of GEM-MACH.

It is important to note that while EnsAI runs significantly faster than GEM-MACH, GEM-MACH produces forecasts for hundreds of chemical fields on 85 vertical levels, while EnsAI only produces a single surface concentration field and does so on a smaller horizontal grid. However, while the wealth of information produced by GEM-MACH can have many uses, the full volume of this information is not always needed for every application. Generating an ensemble of ammonia concentrations at or near the surface to be used for emissions inversions is one example where not having as much information generated is not necessarily a draw back, but where the long run times of the physically-based model may limit how often these inversions can be performed.

# 7 Static Error Covariances

To provide a point of reference for the differences between the GEM-MACH and EnsAI ensemble covariances, we also examine the differences between the GEM-MACH ensemble error covariance and the model error covariance derived in a similar manner to those used operationally at ECCC for surface pollutants (Ménard and Deshaies-Jacques, 2018c; Ménard et al., 2020). ECCC's operational air quality system produces hourly surface analyses of  $O_3$ , NO,  $NO_2$ ,  $SO_2$ ,  $PM_{2.5}$ ,  $PM_{10}$  using error statistics that have a diurnal variation, but are otherwise constant within a two-month period. We refer to these flow-independent error statistics as the 'static' error covariances, although this is a bit of a misnomer as they do contain diurnal patterns. Including these comparisons will help quantify the effect of using flow-dependent error covariances that mirror what is currently used in the operational air quality system at ECCC.

While the ECCC operational error statistics for surface pollutants are formed using the time-average departures of a single forecast from its time-mean and adjusting the error variances using innovation statistics, to isolate the effects of including flow-dependent error covariances, the static error covariances used for this study were formed simply by removing the time-dependence of the GEM-MACH ensemble covariance (but retaining the mean diurnal variation as in the operational system).

# 8 Direct Comparisons of the Covariances

In this section, we describe the methods that will be used for direct comparison between covariances. This will be followed by a comparison of inversions results in Section 9.

Although EnsAI was trained on a 60 member ensemble, it is possible for EnsAI to generate ensembles with more than 60 members. While the speed at which EnsAI can run may allow for much larger ensemble sizes, new strategies for memory management may be required for very large ensembles. In any case, to simplify the comparison between error covariances, for this study the results presented for EnsAI will be for an ensemble of the same size as the original GEM-MACH ensemble.

Uncertainties on quantities derived from the GEM-MACH ensemble due to the limited size of the ensemble will be estimated via bootstrapping, in which the GEM-MACH ensemble members were resampled (with replacement) 50 times. In all line plots shown below, 1–sigma confidence intervals are illustrated by the lightly-shaded region around

the curve. We note that for many of the curves describing the ensemble variance, uncertainties are small enough that the confidence regions are not easily distinguished from the main curve.

#### 8.1 Covariance Parametrization

For the direct comparison between covariances, each covariance is decomposed into variance and correlation components. The error variances can be constructed from the emissions standard deviations  $\sigma_e$  and the surface concentrations standard deviations  $\sigma_c$ , while the correlations are comprised by the univariate emissions correlations  $\rho_{ee}$ , the univariate surface concentration  $\rho_{cc}$ , and the emissions/surface concentration cross-correlations  $\rho_{ec}$ . As EnsAI uses the same emissions ensemble { $\delta e_i$ } as the GEM-MACH ensemble, no comparisons of  $\sigma_e$  or  $\rho_{ee}$  are necessary.

While the ensemble standard deviations  $\sigma_c$  can be represented by a vector and is straightforward to display its values (for example on top of a 2D map), it is more difficult to display the correlations  $\rho_{cc}$  and  $\rho_{ec}$  in full, as these represent matrices, although a single row/column can easily be shown. To aid in summarizing the information in the correlation matrices, we introduce the isotropic correlation length  $L^{iso}$ , which we define at location  $\mathbf{x}$  as

$$L^{\rm iso}(\mathbf{x}) = \sqrt{\frac{2}{\pi}} \int \varrho(\mathbf{x}, \mathbf{x}') |\mathbf{x}'| d\mathbf{x}', \tag{6}$$

where  $\rho$  is the normalize spatial correlation function defined by

$$\varrho(\mathbf{x}, \mathbf{x}') \equiv \frac{\rho(\mathbf{x}, \mathbf{x}')}{\int \rho(\mathbf{x}, \mathbf{x}') d\mathbf{x}'}.$$
(7)

Defined in this manner, the isotropic correlation length is the expectation value of distance taken with respect to  $\rho$ . The factor of  $\sqrt{2/\pi}$  in this definition is included so that for a 2D isotropic Gaussian correlation model,  $L^{\text{iso}}$  coincides with the Daley length-scale (Daley, 1993) (as well as the standard deviation of the Gaussian model). We can extend this definition to anisotropic modes m of the correlation length by introducing a multipole expansion of the correlation length as

$$L[m](\mathbf{x}) \equiv \sqrt{\frac{2}{\pi}} \int e^{im\phi'} \varrho(\mathbf{x}, \mathbf{x}') |\mathbf{x}'| d\mathbf{x}', \qquad (8)$$

where  $\phi'$  is the angle between  $\mathbf{x}'$  and the x-axis. With this definition, we have  $L[0] = L^{\text{iso}}$ , L[1] is proportional to the dipole moments of the function  $\rho$ , and L[2] is related to its quadrupole moments. More details of this expansion can be found in Appendix C.

To better understand this parameterization of the correlation length, example correlations are given in Figure 4, with the values for L[1] and L[2] displayed in the upper left corner of each panel. In these plots, the correlation length modes are taken with respect to the origin. The isotropic correlation length L[0] is the same in each panel and the units of L[m] are set such that L[0] is equal to one. In the top row, all correlations have L[1] = 0 but have varying values for L[2], which results in anisoptric correlations centered on the origin and its major axis pointing in different directions. The correlations in the bottom row of Fig. 4 have L[2] values with the same magnitude as those in the top row, but also have nonzero values for L[1], which skews the correlations in a particular direction.



Figure 4: Examples of correlations and their corresponding L[1] and L[2] values, which are displayed in the upper left corner of each panel. The correlation length modes are taken with respect to the origin. The values for L[0] is the same for all panels and units are selected so that L[0] = 1.

The set of covariance parameters to compare is then  $\{\sigma_c, L_{ec}[m], L_{cc}[m]\}\)$ , where  $L_{ec}[m]$  and  $L_{cc}[m]$  are the correlation length modes for the emissions/surface concentration cross-correlation and surface concentration univariate correlations, respectively. Lastly, since the comparisons of  $L_{ec}[m]$  and  $L_{cc}[m]$  yielded fairly similar conclusions, our analysis will focus on  $L_{ec}[m]$  for brevity and only some comparisons for  $L_{cc}[m]$  will be displayed.

#### 8.2 Comparison Statistics

For each covariance parameter  $\theta \in \{\sigma_c, L_{ec}, L_{cc}\}$ , we will examine the root mean square error  $RMSE(\theta)$  and either the correlation coefficient  $r(\theta)$  or the anomaly correlation coefficient  $ACC(\theta)$ , defined by

$$RMSE(\theta) = \sqrt{\frac{1}{N} \sum_{i}^{N} |\theta_i^{\text{EnsAI}} - \theta_i^{\text{GM}}|^2},$$
(9a)

$$r(\theta) = \frac{\sum_{i}^{N} \left(\theta_{i}^{\mathrm{GM}} - \bar{\theta}^{\mathrm{GM}}\right) \left(\theta_{i}^{\mathrm{EnsAI}} - \bar{\theta}^{\mathrm{EnsAI}}\right)}{\sqrt{\sum_{i}^{N} |\theta_{i}^{\mathrm{GM}} - \bar{\theta}^{\mathrm{GM}}|^{2}} \sqrt{\sum_{i}^{N} |\theta_{i}^{\mathrm{EnsAI}} - \bar{\theta}^{\mathrm{EnsAI}}|^{2}}},$$
(9b)

$$ACC(\theta) = \frac{\sum_{i}^{N} \left(\theta_{i}^{\mathrm{GM}} - \theta_{i}^{\mathrm{c}} - \bar{\theta}^{\mathrm{GM}} + \bar{\theta}^{\mathrm{c}}\right) \left(\theta_{i}^{\mathrm{EnsAI}} - \theta_{i}^{\mathrm{c}} - \bar{\theta}^{\mathrm{EnsAI}} + \bar{\theta}^{\mathrm{c}}\right)}{\sqrt{\sum_{i}^{N} |\theta_{i}^{\mathrm{GM}} - \theta_{i}^{\mathrm{c}} - \bar{\theta}^{\mathrm{GM}} + \bar{\theta}^{\mathrm{c}}|^{2}} \sqrt{\sum_{i}^{N} |\theta_{i}^{\mathrm{EnsAI}} - \theta_{i}^{\mathrm{c}} - \bar{\theta}^{\mathrm{EnsAI}} + \bar{\theta}^{\mathrm{c}}|^{2}}}.$$
(9c)

In the equations above, the superscripts 'EnsAI' and 'GM' denote ensemble quantities derived from the EnsAI and GEM-MACH ensembles, respectively, i indexes each quantity (in space and/or time) with a total of N elements, and overbars denote the mean of that quantity. The same statistics will also be computed for the static error covariances to use as a point of reference, which are computed by replacing all variables that have the 'EnsAI' superscript in Eqs. (9) by the analogous quantities for the static error covariances, which will be labeled with the 'static' superscript. In Eq. (9c), the superscript 'c' refers to a reference value (traditionally a climatological value), which for this work we take as the average value over the whole time period considered (May to August 2016).

#### 8.3 Comparison Results for Error Variances

We begin with a comparison of the standard deviation of the ammonia surface concentrations  $\sigma_c$ . The error standard deviations as a function of time at the five locations shown in Fig. 1 can be seen in the top panels of Figures 5 to 9. These locations correspond to the places with the highest values in the annual emissions inventory and are ordered with decreasing emissions. From these time series, we can see that EnsAI captures a considerable amount of the temporal variation of the GEM-MACH ensemble variances. There are some notable differences, such as EnsAI's over-prediction of variance at the location in Iowa in June (Fig. 6). Another significant difference can be seen in Pennsylvania around June 2-3 (Fig. 8), where EnsAI's variance suddenly spikes. Summary statistics for these time

series are shown in tables to the right of each plot. From these tables, we can see that the correlation coefficient  $r(\sigma_c)$  is significantly higher and  $RMSE(\sigma_c)$  is lower for the EnsAI ensemble as compared to the static error covariance for all locations with the exception of Georgia. For the location in Georgia, both  $r(\sigma_c)$  and  $RMSE(\sigma_c)$  are slightly lower for the static error covariance, although the difference in  $RMSE(\sigma_c)$  has low statistical significance.



Figure 5: Time series of components of the error covariances for the location in North Carolina (exact location is shown in Fig. 1) for the GEM-MACH ensemble (blue), the EnsAI ensemble (orange), and the static covariance (green). The top plot shows the time series of the standard deviation of surface ammonia  $\sigma_c$ . Lower plots show time series for the horizontal correlation length modes  $L_{ec}[m]$  for wavenumbers  $m \leq 2$  for the emissions/surface concentration cross-correlation. Time-mean statistics over the combined time period of July 1-7 and August 1-7 are shown in tables to the right of each plot, where RMSE values in the first and second tables from the top have units of ppbv and km, respectively, and all other table values are unitless.



Figure 6: Same as Fig. 5, but for the location in Iowa.



Figure 7: Same as Fig. 5, but for the location in Ohio.



Figure 8: Same as Fig. 5, but for the location in Pennsylvania.



Figure 9: Same as Fig. 5, but for the location in Georgia.

The largest  $RMSE(\sigma_c)$  values in the model domain generally occur at or near the locations labeled in Fig. 1 (see Figure S2 of the Supplement for  $RMSE(\sigma_c)$  values for whole model domain). The most notable difference between

	ACC		RMSE	
	EnsAI	static	EnsAI	static
$\sigma_c$	$0.73\pm0.01$	$0.36\pm0.01$	$0.36 \pm 0.01$ ppbv	$0.44 \pm 0.01 \text{ ppbv}$
$L_{ec}[0]$	$0.58\pm0.02$	$0.06\pm0.02$	$6.51\pm0.04~\mathrm{km}$	$9.29\pm0.03~\mathrm{km}$
$L_{ec}[1]/L_{ec}[0]$	$0.87\pm0.02$	$0.05\pm0.01$	$0.167 \pm 0.001$	$0.346 \pm 0.001$
$L_{ec}[2]/L_{cc}[0]$	$0.81\pm0.03$	$0.05\pm0.01$	$0.145 \pm 0.001$	$0.253 \pm 0.001$
$L_{cc}[0]$	$0.58\pm0.02$	$0.17\pm0.02$	$5.7\pm0.3~\mathrm{km}$	$8.3\pm0.1~{ m km}$
$L_{cc}[1]/L_{cc}[0]$	$0.58\pm0.02$	$0.06\pm0.01$	$0.182 \pm 0.004$	$0.229 \pm 0.002$
$L_{cc}[2]/L_{cc}[0]$	$0.80\pm0.01$	$0.02\pm0.01$	$0.174 \pm 0.003$	$0.293 \pm 0.002$

Table 1: Domain-wide statistics for components of the error covariances for the combined time period of July 1-7 and August 1-7.

 $RMSE(\sigma_c^{\text{EnsAI}})$  and  $RMSE(\sigma_c^{\text{static}})$  occurs in North Carolina, where at some locations  $RMSE(\sigma_c^{\text{EnsAI}})$  is 3.4 ppbv less than  $RMSE(\sigma_c^{\text{static}})$ .

The values for  $RMSE(\sigma_c)$  and  $ACC(\sigma_c)$  taken over the whole model domain and evaluation time period are shown in the top row of Table 1, from which we can see that  $ACC(\sigma_c)$  is significantly larger for the EnsAI ensemble as compared to the static covariance, indicating that the EnsAI ensemble was much better at capturing the time variation of  $\sigma_c^{\text{GM}}$ . The EnsAI ensemble also has a lower  $RMSE(\sigma_c)$  value than the static ensemble.

#### 8.4 Comparison Results for Spatial Correlations

As an illustration of the time-dependency of the spatial correlations, Figure 10 shows the emissions/surface concentration spatial correlation  $\rho_{ec}$  at the location in North Carolina at three different times, separated by two days and all at 18:00 UTC. Looking at the correlations of the GEM-MACH ensemble in the left column of Fig. 10, we can see that the correlations are often significantly anisotropic and spread out in different directions for different days, reflecting the different meteorological conditions at each time. Comparing this to the correlation in the EnsAI ensemble in the middle column of this figure, it is evident that EnsAI can reasonably reproduce this time-dependence, as the shape of the spatial correlations between the two ensembles look similar. In contrast, the static correlations displayed in the right column of the figure are much closer to being isotropic.



Figure 10: Emissions/surface concentration spatial correlations  $\rho_{ec}$  for the GEM-MACH ensemble (left column), the EnsAI ensemble (middle column), and the static covariance (right column) at the location in North Carolina at three different times.

As described in Section 8.1, to make the examination of the correlation matrix more tractable, we will examine the different modes m of the correlation length  $L_{ec}[m]$  instead of dealing with the full correlation matrix directly. To illustrate the expansion into modes, Figure 11 shows the correlation length modes  $L_{ec}[m]$  as a function of wavenumbers m for the correlations in the top row of Fig. 10. In this example, the isotropic correlation lengths  $L_{ec}[0]$  for the different covariances are between 65 and 90 km, while the anisotropic correlation length modes decrease rapidly with increasing wavenumber. Looking at the absolute value of the correlation length modes (top panel of Fig. 11), we can see that the non-isotropic modes for the static covariance is much smaller than that for the GEM-MACH and EnsAI ensembles, indicating less anisotropy in the static correlation. It is also evident from the bottom two panels in this figure that for most modes  $L_{ec}[m]$  from the EnsAI ensemble is closer to the GEM-MACH values than the static covariance is. As  $L_{ec}[m]$  decreases rapidly with increasing wavenumber, for the remainder of this paper we concentrate mainly on results for modes with  $m \leq 2$  for brevity.



Figure 11: Horizontal correlation length modes  $L_{ec}[m]$  as a function of wavenumber m for the emissions/surface concentration cross-correlation at the location in North Carolina at 2016-06-01 18:00. The top, middle, and bottom panels show the absolute, real, and imaginary parts, respectively.

Time series of  $L_{ec}[m]$  for  $m \leq 2$  for the five locations identified in Fig. 1 are shown in the bottom four panels of Figs. 5 to 9. In all of these plots,  $L_{ec}[m]$  for the EnsAI ensemble are much closer to the GEM-MACH ensemble value than the static covariance is and captures most of its time variation. While a diurnal pattern is evident in the GEM-MACH ensemble, there is significant non-diural time variation throughout the week that is dependent on meteorological conditions. In Fig. 5, we can see that the EnsAI correlation length modes generally follow the pattern in the original GEM-MACH ensemble, but the results are even more striking for the other four locations (Figs. 6 to 9), where the differences between the EnsAI and GEM-MACH ensembles are often smaller (or not much larger) than the sample variation in the GEM-MACH ensemble, or in other words, the differences between these values are statistically insignificant throughout much of the evaluation time period. Looking at the evaluation statistics taken over the full test time period, which are displayed in the tables on the right side of Figs. 5 to 9, the *RMSE* values are smaller and correlation coefficients larger for the EnsAI ensemble than for the static covariance for all locations. Many correlation coefficients values for the EnsAI ensemble are above 0.8, while those for the static covariance are mostly below 0.3.

The RMSE values of  $L_{ec}[m]$  for the EnsAI ensemble are generally much smaller than those for the static covariance at most grid points (see Figure S3 of the Supplement). In many parts of the model domain the EnsAI ensemble has values  $RMSE(L_{ec}[0]) \approx 10$  km,  $RMSE(L_{ec}[1]/L_{ec}[0]) \approx 0.2$ , and  $RMSE(L_{ec}[2]/L_{ec}[0]) \approx 0.1$ , while at the same grid points the static covariance has values of  $RMSE(L_{ec}[0]) \approx 20$  km,  $RMSE(L_{ec}[1]/L_{ec}[0]) \approx 0.6$ , and  $RMSE(L_{ec}[2]/L_{ec}[0]) \approx 0.5$ . For reference, the mean value of  $L_{ec}[0]$  over the model domain is roughly equal to 80 km for the GEM-MACH ensemble.

The ACC and RMSE values for both  $L_{ec}[m]$  and  $L_{cc}[m]$ , taken over the whole model domain and evaluation time period, are displayed in the bottom six rows of Table 1. The RMSE values in Table 1 for the EnsAI ensemble are about 3 km, 0.2, and 0.1 lower then those for the static covariance for  $L_{ec}[0]$ ,  $L_{ec}[1]/L_{ec}[0]$ , and  $L_{ec}[2]/L_{ec}[0]$ ,

respectively. The differences in RMSE values for  $L_{cc}[m]$  are similar, with the exception of  $L_{cc}[1]/L_{cc}[0]$  which is only about 0.05 lower for the EnsAI ensemble. The ACC values for the EnsAI ensemble is significantly higher than for the static covariance for all correlation length modes given in Table 1. The ACC values for the EnsAI ensemble are all between 0.58 and 0.87, while the ACC values for the static covariance are all lower than 0.2. From this we can conclude that the EnsAI ensemble can accurately reproduce the time variations in the horizontal correlations that are not captured by the diurnal pattern of the static covariances.

To get an intuitive understanding of what these differences in correlation length modes correspond to in terms of spatial patterns, Figure 12 shows another set of example correlations that illustrate the deviations in correlation length mode values mentioned in the previous paragraphs. The leftmost panel of Fig. 12 shows an isotropic correlation with L[0] = 80 km. For illustrative purposes, we will think of this isotropic correlation as a target correlation (i.e. from the GEM-MACH ensemble) and compare it to correlations with different L[m] values to see how different the spatial patterns are. The correlation shown in the panel second from the left has L[0], L[1], and L[2] values that deviate from the isotropic correlation in the leftmost panel by amounts equal to the RMSE values found in Table 1 for EnsAI. This correlation is only slightly anisotropic and slightly smaller than the correlation in the leftmost panel. The correlation in the right has values that deviate from the isotropic correlation equal to the RMSE values for the static covariance in Table 1 and is noticeably anisotropic. The correlation in the rightmost panel has values that deviate from the isotropic correlation by amounts equal to the difference between EnsAI and static covariance values that are commonly found in the model domain (Figure S3 of the Supplement), which produces a correlation with a drastically different spatial pattern than the correlation in the leftmost panel.



Figure 12: Example correlations with their corresponding L[0], L[1], and L[2] values displayed at the top of each panel.

To conclude this section, we briefly examine the higher correlation length modes with m > 2. Figure 13 shows the RMSE and ACC values of  $L_{ec}[m]/L_{ec}[0]$ , taken over the evaluation time period and model domain, for m between 1 and 15. From the top plot, we can see that the RMSE values for both the EnsAI ensemble and static covariance decrease with increasing wavenumber, as all three covariances typically have less features on smaller scales (the decreasing values of  $L_{ec}[m]$  with increasing m that was shown in the example in Fig. 11 is a typical feature of the horizontal correlations). The RMSE values for the EnsAI ensemble are smaller than those for the static covariance for all wavenumbers, but the difference between the two decrease with increasing values of m and nearly converge by m = 15. As seen in bottom panel of Fig. 13, the ACC values for the EnsAI ensemble is above 0.85 for m = 1 and decreases with increasing m, but is still at a value just below 0.5 by m = 15. In contrast, the ACC values for the static covariance show little change with wavenumber and is below 0.1 for all wavenumbers.



Figure 13: RMSE and ACC values for the horizontal correlation length modes for the emissions/surface concentration correlation  $L_{ec}[m]$  as a function of wavenumber m for the EnsAI ensemble (orange) and the static covariance (green).

### 9 Inversions

Sections 8.3 and 8.4 showed the results of direct comparisons between the covariances to determine how well EnsAI replicated the GEM-MACH ensemble perturbations. For some applications, such as model uncertainty quantification, the determination of the model's error covariance is the end goal. However, for inversions/data assimilation, the objective is to produce an analysis (as well as forecasts initialized by the analysis in the case of data assimilation). In this section, we compare emission inversion results made using the different model error covariances.

As the inversions are only sensitive to the model error covariances at times and locations near the available observations, not all differences between the error covariances will show in the comparison between the increments and will be dependent on the particular observations used. As such, pairing the comparison of inversion results with the direct comparison between the covariances in the previous two sections provides a robust examination between the different error covariances.

The ammonia emission increments produced using the inversion scheme described in Section 2 for the two test time periods are displayed in Figure 14. The left column shows the emission increments produced by inversions that used the original GEM-MACH ensemble. As found in Sitwell et al. (2022), the inversions increase the ammonia emissions in most parts of the model domain, with a few decreases in June in some parts of North Carolina and in the midwestern United States. We note that the emissions increments are smaller than those computed in Sitwell et al. (2022), since the inversions in this work each only use a week's worth of observations instead of a full month's worth. However, since our goal is only to determine how well inversions that use EnsAI reproduces the increments made using the original GEM-MACH ensemble, this should not be relevant for our comparisons.

The middle column in Fig. 14 shows the difference between the increments produced using the EnsAI ensemble and those produced using the GEM-MACH ensemble (as displayed in the left column), while the right column shows the difference between the increments made using the static covariance and those made using the GEM-MACH ensemble. Looking at the middle column of Fig. 14, the EnsAI and GEM-MACH increments are mostly within 5 mg s<sup>-1</sup> km<sup>-2</sup> of one another, corresponding to being within roughly 15% of one another. The only major exception to this is in

north-wester Iowa in the June inversion, where the EnsAI increment is only around -12 mg s<sup>-1</sup> km<sup>-2</sup> instead of about -25 mg s<sup>-1</sup> km<sup>-2</sup> for the GEM-MACH inversion. By contrast, the right column of Fig. 14 shows that there are a number of locations (North Carolina, Pennsylvania, southern Nebraska, central Texas) where the static covariance and GEM-MACH increments differ by over 15 mg s<sup>-1</sup> km<sup>-2</sup>. The increments produced using the static covariances differs from the GEM-MACH increment by a larger amount than the EnsAI increment does in most locations within the model domain.



Figure 14: Ammonia emissions increments for inversions using CrIS ammonia retrievals for the week of 2016-06-01 to 2016-06-07 (top row) and for the week of 2016-08-01 to 2016-08-07 (bottom row). The left column shows the increments produced with inversions that use the GEM-MACH ensemble. The middle column shows the difference between increments produced using the EnsAI ensemble and increments produced using the GEM-MACH ensemble, while the right column shows the difference between the static error covariance increments and the GEM-MACH ensemble increments.

# 10 Conclusions

Obtaining flow-dependent model error covariances is an important part of modern data assimilation and inversion systems. Ensemble-based methods have become a popular choice for adding flow-dependence to the error covariances as it does not require the use of a model adjoint that can be time consuming to create and maintain. The downside however is the need to repeatedly run the model to generate the ensemble, which can be a significant burden on computational resources for atmospheric models. In this work we introduced EnsAI, an AI-based ensemble generation model for atmospheric chemical constituents that can run many orders of magnitude faster than traditional physics-based models.

When trained on an existing ensemble of atmospheric ammonia concentrations generated using the GEM-MACH model, EnsAI can accurately reproduce many features of the original ensemble. By comparing the original GEM-MACH and EnsAI ensembles to a static error covariance, it was shown that including flow-dependence in the error covariances has a considerable effect on the covariances and that EnsAI can encode much of the meteorological-dependent features of the original ensemble. Furthermore, it was demonstrated that including only surface information

in EnsAI was sufficient to reproduce most of the time-dependent features of the ensemble, thereby allowing EnsAI to produce accurate ensemble information for ammonia with much smaller memory requirements as compare to systems that require multiple vertical levels.

Many geophysical AI-based models use extensive training data sets, such as the ERA5 or CAMS analysis/reanalysis. However, EnsAI was trained on a much smaller data set that was only comprised of 15 weeks of data. This suggests that ensemble generation systems that are only designed to predict model perturbations (instead of the model trajectory itself) seems to produce good results using much smaller training data sets.

An ensemble-based emissions inversion system for ammonia was used as a test application for EnsAI. Accurate emissions/concentration spatial correlations are necessary for proper spatial source attribution of observations in emissions inversions. As such, capturing the time-dependency of these correlations is important for successful inversions, which can be done by using ensembles in the inversion system. It was demonstrated that EnsAI-based emissions inversions produced very similar results to inversions performed using the GEM-MACH ensemble. Being able to generate ensembles in a fraction of the time required for physics-based models significantly eases the computational cost of running such inversion systems.

While application to emissions inversion systems is a substantial development by itself, the longer term goal is expand EnsAI to more atmospheric chemical constituents to use in chemical data assimilation. Although EnsAI was able to reproduce most the time-dependent features of the ammonia ensemble with only surface fields as input, more vertical levels may be needed if applied to other species with longer atmospheric life times.

# Acknowledgments

The author would like to thank Colin Lee for discussions on AI and GPU computing, Mark Shephard for providing advice on utilizing the CrIS ammonia retreivals, NOAA Comprehensive Large Array-data Stewardship System (CLASS; Liu et al. (2014)) for providing the CrIS Level 1 and Level 2 CrIS REDRO and NUCAPS input atmospheric state data, and the GEM-MACH development team for the creation and maintenance of the GEM-MACH model. The author would also like to thank Colin Lee, Dominique Brunet, Vikram Khade, and Mark Buehner for providing feedback on a draft of this manuscript.

### Appendix

### A Ensemble-Variational Emissions Inversion

The emissions inversions were performed using a variational algorithm, in which a cost function is minimized to find the optimal combination between a priori (background) and observational data. The cost function is minimized with respect to the increment  $\Delta \mathbf{x}$ , which when added to the background state  $\mathbf{x}^{b}$  yields the analysis  $\mathbf{x}^{a} = \mathbf{x}^{b} + \Delta \mathbf{x}$ . In our case, the cost function J is given by

$$J = \frac{1}{2} \Delta \mathbf{x}^{\mathrm{T}} \mathbf{B}^{-1} \Delta \mathbf{x} + \frac{1}{2} (\mathbf{d} - \mathbf{H} \Delta \mathbf{x})^{\mathrm{T}} \mathbf{R}^{-1} (\mathbf{d} - \mathbf{H} \Delta \mathbf{x}),$$
(10)

where  $\mathbf{d} = \mathbf{y} - H(\mathbf{x}^{b})$  is the difference between the observations  $\mathbf{y}$  and the background in observation space. *H* is the (nonlinear) observation operator that maps the model state into observation space, while **H** denotes its linearization. The first term in the right hand side of Equation (10) quantifies the deviation from the background state and is weighted by the inverse of the background error covariance **B**. Similarly, the second term in the right hand side of Eq. (10) quantifies the deviation from the observation error covariance **R**.

In the case of an emissions inversion where the unobserved emissions  $\mathbf{e}$  are constrained using observations of the atmospheric concentration  $\mathbf{c}$ , using Eq. (2), the cost function becomes

$$J = \frac{1}{2} \Delta \mathbf{e}^{\mathrm{T}} \mathbf{B}_{ee}^{-1} \Delta \mathbf{e} + \frac{1}{2} (\mathbf{d} - \mathbf{H} \Delta \mathbf{c})^{\mathrm{T}} \mathbf{R}^{-1} (\mathbf{d} - \mathbf{H} \Delta \mathbf{c}),$$
(11)

where  $\Delta \mathbf{e}$  and  $\Delta \mathbf{c}$  are the emissions and atmospheric concentration increments, respectively. Since we assume that uncertainties in the atmospheric concentrations are due solely to the uncertainties in emissions, the emissions and concentration increments are related to each other by

$$\Delta \mathbf{e} = \mathbf{B}_{ec} \mathbf{B}_{cc}^{-1} \Delta \mathbf{c}. \tag{12}$$

It is useful to rewrite the cost function in terms of a control vector  $\chi$ , which defines the control vector transforms

$$\Delta \mathbf{c} = \mathbf{B}_{cc}^{1/2} \boldsymbol{\chi},\tag{13a}$$

$$\Delta \mathbf{e} = \mathbf{B}_{ee}^{1/2} \boldsymbol{\chi},\tag{13b}$$

where  $\mathbf{B}_{cc}^{1/2}$  and  $\mathbf{B}_{ee}^{1/2}$  are the square roots of  $\mathbf{B}_{cc}$  and  $\mathbf{B}_{ee}$ , respectively. The cost function can then be expressed as

$$J(\boldsymbol{\chi}) = \frac{1}{2} \boldsymbol{\chi}^{\mathrm{T}} \boldsymbol{\chi} + \frac{1}{2} (\mathbf{d} - \mathbf{H} \mathbf{B}_{cc}^{1/2} \boldsymbol{\chi})^{\mathrm{T}} \mathbf{R}^{-1} (\mathbf{d} - \mathbf{H} \mathbf{B}_{cc}^{1/2} \boldsymbol{\chi}).$$
(14)

The emission inversion then consists of minimizing this cost function with respect to the control vector  $\chi$ , followed by the transformation of the optimal control vector to the emissions (or concentration) increment given by Equation (13).

Ensemble-variational methods are a class of variational algorithms that model the background error covariance using an ensemble and is the inversion method used in this study. As mentioned in Section 2.3, covariances that are formed solely from ensemble pertrubations will in general have spurious correlations and low rank due to the limited size of the ensemble. As such, localization of the ensemble covariances is included, as shown in Eq. (3). With this,  $\mathbf{B}_{cc}^{1/2}$  and  $\mathbf{B}_{ee}^{1/2}$  can then be expressed as

$$\mathbf{B}_{cc}^{1/2} = \frac{1}{\sqrt{N-1}} \begin{bmatrix} \operatorname{diag}(\delta \mathbf{c}_1) \mathcal{L}^{1/2} & \dots & \operatorname{diag}(\delta \mathbf{c}_N) \mathcal{L}^{1/2} \end{bmatrix},$$
(15a)

$$\mathbf{B}_{ee}^{1/2} = \frac{1}{\sqrt{N-1}} \begin{bmatrix} \operatorname{diag}(\delta \mathbf{e}_1) \mathcal{L}^{1/2} & \dots & \operatorname{diag}(\delta \mathbf{e}_N) \mathcal{L}^{1/2} \end{bmatrix},$$
(15b)

where  $\{\delta \mathbf{e}_i\}$  and  $\{\delta \mathbf{c}_i\}$  are emissions and concentration perturbation ensembles, respectively,  $\operatorname{diag}(\cdot)$  is the diagonalization operator, and  $\mathcal{L}^{1/2}$  is the square root of the localization matrix. Here, we have chosen the same localization matrix for both the emissions and concentrations, which was defined to be isotropic with a half-width at half-maximum of 120 km with the shape prescribed by the fifth-order function of Gaspari and Cohn (1999).

### **B** Components of the U-Net

The different components that comprise the U-Net used for EnsAI are described in this appendix. In this appendix, all references to graphical representations of the U-Net components refer to Fig. 3.

Residual blocks (He et al., 2016) (represented by the blue right arrows) are used throughout the network and consist of two sequential constitutional layers with a residual connection. Each constitutional layer is comprised by (in sequential order) a  $3 \times 3$  two-dimensional convolution, a group normalization (Wu and He, 2018), and a Sigmoid Linear Unit (SiLU) (Hendrycks and Gimpel, 2016) activation function.

In the encoding branch, the residual blocks are used to increase the number of channels, while a  $2 \times 2$  max pooling (represented by the orange downward arrows) is used to half each spatial dimension. In the decoding branch, upscaling blocks (denoted by the green upward arrows), consisting of a nearest neighbor upscaling followed by a  $3 \times 3$  two-dimensional convolution, are used to double each spatial dimension and half the number of channels. Skip connections (represented by the grey dashed arrows), comprised of a copy and concatenation, are made across the encoding and decoding branches. A final two-dimensional convolution is made at the end of the decoding branch (denoted by the purple right arrow) to reduce the array to a single channel that represents the surface ammonia concentration perturbation.

## C Details of Correlation Length Decomposition

With the decomposition of the correlation length defined in Equation (8), the first three modes are given by

$$L[0] = \sqrt{\frac{2}{\pi}} \int r^2 \varrho(r,\theta) dr d\theta, \qquad (16a)$$

$$L[1] = \sqrt{\frac{2}{\pi}} \int (\cos(\phi) + i\sin(\phi))r^2 \varrho(r,\theta) dr d\theta$$
  
=  $\sqrt{\frac{2}{\pi}} \int (x + iy)\varrho(x,y) dx dy,$  (16b)

$$\begin{split} L[2] &= \sqrt{\frac{2}{\pi}} \int (\cos(2\phi) + i\sin(2\phi))r^2 \varrho(r,\theta) dr d\theta \\ &= \sqrt{\frac{2}{\pi}} \int (\cos^2(\phi) - \sin^2(\phi) + i\sin(\phi)\cos(\phi))r^2 \varrho(r,\theta) dr d\theta \\ &= \sqrt{\frac{2}{\pi}} \int (x^2 - y^2 + ixy) \frac{\varrho(x,y)}{\sqrt{x^2 + y^2}} dx dy, \end{split}$$
(16c)

where in these equations the correlations are computed between the origin and the point (x, y), which when expressed in polar coordinates is a distance r from the origin at an angle  $\theta$  from the x-axis.

Examining the the integral in Equation (16a), at each angle  $\theta$ , the integral over r is equal to the (one-sided) variance of  $\rho$  in that particular direction. The integral over  $\theta$  then takes the angle-averaged value of this one-sided spatial variance. As an illustration of this quantity, consider the spatial correlation function of a 2D isotropic Gaussian  $\rho_G$ , given by

$$\rho_G(x,y) = \exp\left(-\frac{x^2 + y^2}{2L_G^2}\right),\tag{17}$$

where  $L_G$  is the scale parameter of the Gaussian function, so that its normalized spatial correlation  $\rho_G$  is

$$\varrho_G(x,y) = \frac{1}{2\pi L_G^2} \exp\left(-\frac{x^2 + y^2}{2L_G^2}\right).$$
(18)

In this case, L[0] is given by

$$L[0] = \frac{1}{\sqrt{2}\pi^{3/2}L_G^2} \int_0^\infty r^2 \exp\left(-\frac{r^2}{2L_G^2}\right) dr \int_0^{2\pi} d\theta$$
  
=  $\frac{\sqrt{2}}{\sqrt{\pi}L_G^2} \int_0^\infty r^2 \exp\left(-\frac{r^2}{2L_G^2}\right) dr$   
=  $L_G$ , (19)

which shows that L[0] is equal to the Gaussian scale parameter  $L_G$  in this example.

From the second line of Equation (16b), it is apparent that the real and imaginary parts of the L[1] are proportional to the dipole moments of  $\rho$  in the x and y directions, respectively. Functions that are symmetric about the x-axis (y-axis) will have vanishing real (imaginary) parts of L[1].

The third line of Equation (16c) shows that L[2] is related to the quadrupole moments of the function  $\rho(x, y)/\sqrt{x^2 + y^2}$ , where the real part is proportional to the difference between the x, x and y, y quadrupole components and the imaginary part is proportional to the x, y component.

### References

- Anna Agustí-Panareda, Jérôme Barré, Sébastien Massart, Antje Inness, Ilse Aben, Melanie Ades, Bianca C Baier, Gianpaolo Balsamo, Tobias Borsdorff, Nicolas Bousserez, Souhail Boussetta, Michael Buchwitz, Luca Cantarello, Cyril Crevoisier, Richard Engelen, Henk Eskes, Johannes Flemming, Sébastien Garrigues, Otto Hasekamp, Vincent Huijnen, Luke Jones, Zak Kipling, Bavo Langerock, Joe McNorton, Nicolas Meilhac, Stefan Noël, Mark Parrington, Vincent-Henri Peuch, Michel Ramonet, Miha Razinger, Maximilian Reuter, Roberto Ribas, Martin Suttie, Colm Sweeney, Jérôme Tarniewicz, and Lianghai Wu. The cams greenhouse gas reanalysis from 2003 to 2020. *Atmospheric Chemistry and Physics*, 23(6):3829–3859, 2023.
- Ross N Bannister. A review of forecast error covariance statistics in atmospheric variational data assimilation. i: Characteristics and measurements of forecast error covariances. *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 134(637):1951– 1970, 2008.
- Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.
- Cristian Bodnar, Wessel P. Bruinsma, Ana Lucic, Megan Stanley, Anna Allen, Johannes Brandstetter, Patrick Garvan, Maik Riechert, Jonathan A. Weyn, Haiyu Dong, Jayesh K. Gupta, Kit Thambiratnam, Alexander T. Archibald, Chun-Chieh Wu, Elizabeth Heider, Max Welling, Richard E. Turner, and Paris Perdikaris. Aurora: A foundation model for the earth system, 2024. URL https://arxiv.org/abs/2405.13063.

- L Bruhwiler, E Dlugokencky, K Masarie, M Ishizawa, A Andrews, J Miller, C Sweeney, P Tans, and D Worthy. Carbontracker-ch4: an assimilation system for estimating emissions of atmospheric methane. *Atmospheric Chemistry and Physics*, 14(16):8269–8293, 2014.
- Mark Buehner, PL Houtekamer, Cecilien Charette, Herschel L Mitchell, and Bin He. Intercomparison of variational data assimilation and the ensemble kalman filter for global deterministic nwp. part i: Description and singleobservation experiments. *Monthly Weather Review*, 138(5):1550–1566, 2010.
- Mark Buehner, Ron McTaggart-Cowan, Alain Beaulne, Cécilien Charette, Louis Garand, Sylvain Heilliette, Ervig Lapalme, Stéphane Laroche, Stephen R Macpherson, Josée Morneau, and Ayrton Zadra. Implementation of deterministic weather forecasting systems based on ensemble-variational data assimilation at environment canada. part i: The global system. *Monthly Weather Review*, 143(7):2532–2559, 2015.
- Mark Buehner, Jean-François Caron, Ervig Lapalme, Alain Caya, Ping Du, Yves Rochon, Sergey Skachko, Maziar Bani Shahabadi, Sylvain Heilliette, Martin Deshaies-Jacques, Weiguang Chang, and Michael Sitwell. The modular and integrated data assimilation system at environment and climate change canada (midas v3. 9.1). *Geoscientific Model Development Discussions*, 2024:1–29, 2024.
- Hansen Cao, Daven K Henze, Mark Shephard, Enrico Dammers, Karen Cady-Pereira, Matthew Alvarado, Chantelle Lonsdale, Gan Luo, Fangqun Yu, Liye Zhu, Camille Danielson, and Eric Edgerton. Inverse modeling of nh3 sources using cris remote sensing measurements. *Environ. Res. Lett.*, 15(10):104082, 2020. doi: https://doi.org/10.1088/ 1748-9326/abb5cc.
- Kang Chen, Tao Han, Junchao Gong, Lei Bai, Fenghua Ling, Jing-Jia Luo, Xi Chen, Leiming Ma, Tianning Zhang, Rui Su, Yuanzheng Ci, Bin Li, Xiaokang Yang, and Wanli Ouyang. Fengwu: Pushing the skillful global medium-range weather forecast beyond 10 days lead. arXiv preprint arXiv:2304.02948, 2023a.
- Kun Chen, Lei Bai, Fenghua Ling, Peng Ye, Tao Chen, Jing-Jia Luo, Hao Chen, Yi Xiao, Kang Chen, Tao Han, and Wanli Ouyang. Towards an end-to-end artificial intelligence driven global weather forecasting system. *arXiv* preprint arXiv:2312.12462, 2023b.
- Lei Chen, Xiaohui Zhong, Feng Zhang, Yuan Cheng, Yinghui Xu, Yuan Qi, and Hao Li. Fuxi: A cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and Atmospheric Science*, 6(1):190, 2023c.
- Jean Côté, Jean-Guy Desmarais, Sylvie Gravel, André Méthot, Alain Patoine, Michel Roch, and Andrew Staniforth. The operational cmc–mrb global environmental multiscale (gem) model. part ii: Results. *Mon. Weather. Rev.*, 126 (6):1397–1418, 1998a. doi: https://doi.org/10.1175/1520-0493(1998)126(1397:TOCMGE)2.0.CO;2.
- Jean Côté, Sylvie Gravel, André Méthot, Alain Patoine, Michel Roch, and Andrew Staniforth. The operational cmc– mrb global environmental multiscale (gem) model. part i: Design considerations and formulation. *Mon. Weather: Rev.*, 126(6):1373–1395, 1998b. doi: https://doi.org/10.1175/1520-0493(1998)126(1373:TOCMGE)2.0.CO;2.
- Roger Daley. Atmospheric data analysis. Number 2. Cambridge university press, 1993.
- Gérald Desroziers, Jean-Thomas Camino, and Loïk Berre. 4denvar: link with 4d state formulation of variational assimilation and different possible implementations. *Quarterly Journal of the Royal Meteorological Society*, 140 (684):2097–2110, 2014.
- Geir Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans*, 99(C5):10143–10162, 1994.
- David Fairbairn, Stephen R Pring, Andrew C Lorenc, and Ian Roulstone. A comparison of 4dvar with ensemble data assimilation methods. *Quarterly Journal of the Royal Meteorological Society*, 140(678):281–294, 2014.
- Sebastien Garrigues, Julien Chimot, Melanie Ades, Antje Inness, Johannes Flemming, Zak Kipling, Istvan Laszlo, Angela Benedetti, Roberto Ribas, Soheila Jafariserajehlou, Bertrand Fougnie, Shobha Kondragunta, Richard Engelen, Vincent-Henri Peuch, Mark Parrington, Nicolas Bousserez, Margarita Vazquez-Navarro, and Anna Agusti-Panareda. Monitoring multiple satellite aerosol optical depth (aod) products within the copernicus atmosphere monitoring service (cams) data assimilation system. *Atmospheric Chemistry and Physics*, 22(22):14657–14692, 2022.
- Gregory Gaspari and Stephen E Cohn. Construction of correlation functions in two and three dimensions. Q. J. Roy. Meteor. Soc., 125(554):723–757, 1999. doi: https://doi.org/10.1002/qj.49712555417.
- Claude Girard, André Plante, Michel Desgagné, Ron McTaggart-Cowan, Jean Côté, Martin Charron, Sylvie Gravel, Vivian Lee, Alain Patoine, Abdessamad Qaddouri, Michel Roch, Lubos Spacek, Monique Tanguay, Paul Vaillancourt, and Ayrton Zadra. Staggered vertical discretization of the canadian environmental multiscale (gem) model using a coordinate of the log-hydrostatic-pressure type. *Mon. Weather. Rev.*, 142(3):1183–1196, 2014. doi: https://doi.org/10.1175/MWR-D-13-00255.1.

- W Gong, PA Makar, J Zhang, J Milbrandt, S Gravel, KL Hayden, AM Macdonald, and WR Leaitch. Modelling aerosol-cloud-meteorology interaction: A case study with a fully coupled air quality model (gem-mach). *Atmos. Environ.*, 115:695–715, 2015. doi: https://doi.org/10.1016/j.atmosenv.2015.05.062.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.

Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415, 2016.

- Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla, Xavier Abellan, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot, Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger, Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux, Cristina Lopez, Philippe abd Lupu, Gabor Radnoti, Patricia de Rosnay, Iryna Rozum, Freja Vamborg, Sebastien Villaume, and Jean-Noël Thépaut. The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. Advances in neural information processing systems, 33:6840–6851, 2020.
- Peter L Houtekamer and Fuqing Zhang. Review of the ensemble kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 144(12):4489–4532, 2016.
- Vincent Huijnen, Kazuyuki Miyazaki, Johannes Flemming, Antje Inness, Takashi Sekiya, and Martin G Schultz. An intercomparison of tropospheric ozone reanalysis products from cams, cams interim, tcr-1, and tcr-2. *Geoscientific* model development, 13(3):1513–1544, 2020.
- Antje Inness, Frank Baier, Angela Benedetti, Idir Bouarar, Simon Chabrillat, Hannah Clark, Cathy Clerbaux, Pierre Coheur, RJ Engelen, Quentin Errera, J. Flemming, M George, C Granier, J Hadji-Lazaro, V Huijnen, D Hurtmans, L Jones, JW Kaiser, J Kapsomenakis, K Lefever, J Leitão, M Razinger, A Richter, MG Schultz, AJ Simmons, M Suttie, O Stein, JN Thépaut, V Thouret, M Vrekoussis, C Zerefos, and the MACC team. The macc reanalysis: an 8 yr data set of atmospheric composition. *Atmospheric chemistry and physics*, 13(8):4073–4109, 2013.
- Antje Inness, Melanie Ades, Anna Agustí-Panareda, Jérôme Barré, Anna Benedictow, Anne-Marlene Blechschmidt, Juan Jose Dominguez, Richard Engelen, Henk Eskes, Johannes Flemming, Vincent Huijnen, Luke Jones, Zak Kipling, Sebastien Massart, Mark Parrington, Vincent-Henri Peuch, Miha Razinger, Samuel Remy, Michael Schulz, and Martin Suttie. The cams reanalysis of atmospheric composition. *Atmospheric Chemistry and Physics*, 19(6): 3515–3556, 2019.
- WA Lahoz, Q Errera, R Swinbank, and D Fonteyn. Data assimilation of stratospheric constituents: a review. Atmospheric Chemistry and Physics, 7(22):5745–5773, 2007.
- Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Oriol Vinyals, Jacklynn Stott, Alexander Pritzel, Shakir Mohamed, and Peter Battaglia. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.
- Lizao Li, Robert Carver, Ignacio Lopez-Gomez, Fei Sha, and John Anderson. Generative emulation of weather forecast ensembles with diffusion models. *Science Advances*, 10(13):eadk4489, 2024a.
- Yonghui Li, Wei Han, Hao Li, Wansuo Duan, Lei Chen, Xiaohui Zhong, Jincheng Wang, Yongzhu Liu, and Xiuyu Sun. Fuxi-en4dvar: An assimilation system based on machine learning weather forecasting model ensuring physical constraints. *Geophysical Research Letters*, 51(22):e2024GL111136, 2024b.
- Q. Liu, W. Wolf, T. Reale, A. Sharma, and NOAA JPSS Program Office. Noaa unique combined atmospheric processing system (nucaps) environmental data record (edr) products, 2014.
- Andrew C Lorenc. The potential of the ensemble kalman filter for nwp—a comparison with 4d-var. *Quarterly Journal* of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography, 129(595):3183–3203, 2003.
- Richard Ménard and Martin Deshaies-Jacques. Evaluation of analysis by cross-validation. part i: Using verification metrics. *Atmosphere*, 9(3):86, 2018a.
- Richard Ménard and Martin Deshaies-Jacques. Evaluation of analysis by cross-validation, part ii: Diagnostic and optimization of analysis error covariance. *Atmosphere*, 9(2):70, 2018b.
- Richard Ménard and Martin Deshaies-Jacques. Error covariance estimation methods based on analysis residuals and its application to air quality surface observation networks. In *Air Pollution Modeling and its Application XXV 35*, pages 249–253. Springer, 2018c.

- Richard Ménard, Jean-François Cossette, and Martin Deshaies-Jacques. On the complementary role of data assimilation and machine learning: An example derived from air quality analysis. In *Computational Science–ICCS 2020:* 20th International Conference, Amsterdam, The Netherlands, June 3–5, 2020, Proceedings, Part VI 20, pages 212– 224. Springer, 2020.
- K Miyazaki, HJ Eskes, and K Sudo. Global nox emission estimates derived from an assimilation of omi tropospheric no2 columns. *Atmospheric Chemistry and Physics*, 12(5):2263–2288, 2012a.
- Kazuyuki Miyazaki, Takashi Maki, Prabir Patra, and Takakiyo Nakazawa. Assessing the impact of satellite, aircraft, and surface observations on co2 flux estimation using an ensemble-based 4-d data assimilation system. *Journal of Geophysical Research: Atmospheres*, 116(D16), 2011.
- Kazuyuki Miyazaki, HJ Eskes, K Sudo, M Takigawa, M Van Weele, and KF Boersma. Simultaneous assimilation of satellite no2, o3, co, and hno3 data for the analysis of tropospheric chemical composition and emissions. *Atmospheric Chemistry and Physics*, 12(20):9545–9579, 2012b.
- Michael Moran, Sylvain Ménard, D. Talbot, P. Huang, Paul Makar, W Gong, Hugo Landry, Sylvie Gravel, S. Gong, L.-P Crevier, A. Kallaur, and Mourad Sassi. Particulate-matter forecasting with gem-mach15, a new canadian air-quality forecast model. *Air Pollution Modelling and Its Application XX*, pages 289–292, 2010.
- Soon-Young Park, Uzzal Kumar Dash, Jinhyeok Yu, Keiya Yumimoto, Itsushi Uno, and Chul Han Song. Implementation of an ensemble kalman filter in the community multiscale air quality model (cmaq model v5.1) for data assimilation of ground-level pm 2.5. *Geoscientific Model Development Discussions*, 2021:1–35, 2021.
- Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, Pedram Hassanzadeh, Karthik Kashinath, and Animashree Anandkumar. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. arXiv preprint arXiv:2202.11214, 2022.
- Radenko Pavlovic, Jack Chen, Kerry Anderson, Michael D Moran, Paul-André Beaulieu, Didier Davignon, and Sophie Cousineau. The firework air quality forecast system with near-real-time biomass burning emissions: Recent developments and evaluation of performance for the 2015 north american wildfire season. J. Air. Waste. Manage., 66(9):819–841, 2016. doi: https://doi.org/10.1080/10962247.2016.1158214.
- Drew C Pendergrass, Daniel J Jacob, Hannah Nesser, Daniel J Varon, Melissa Sulprizio, Kazuyuki Miyazaki, and Kevin W Bowman. Cheereio 1.0: a versatile and user-friendly ensemble-based chemical data assimilation and emissions inversion platform for the geos-chem chemical transport model. *Geoscientific Model Development*, 16 (16):4793–4810, 2023.
- W Peters, MC Krol, GR Van Der Werf, S Houweling, CD Jones, J Hughes, K Schaefer, KA Masarie, AR Jacobson, JB Miller, CH Cho, M RAMONET, M SCHMIDT, L CIATTAGLIA, F APADULA, D HELTAI, F MEINHARDT, AG DI SARRA, S PIACENTINO, D SFERLAZZO, T AALTO, J HATAKKA, J STRöM, L HASZPRA, HAJ MEIJER, S VAN DER LAAN, REM NEUBERT, A JORDAN, X RODó, JA MORGUÍ, AT VERMEULEN, E POPA, K ROZANSKI, M ZIMNOCH, AC MANNING, M LEUENBERGER, C UGLIETTI, AJ DOLMAN, P CIAIS, M HEIMANN, and PP TANS. Seven years of recent european net terrestrial carbon dioxide exchange constrained by atmospheric observations. *Global Change Biology*, 16(4):1317–1337, 2010.
- Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, Remi Lam, and Matthew Willson. Gencast: Diffusionbased ensemble forecasting for medium-range weather. arXiv preprint arXiv:2312.15796, 2023.
- Florence Rabier, Heikki Järvinen, E Klinker, J-F Mahfouf, and A Simmons. The ecmwf operational implementation of four-dimensional variational assimilation. i: Experimental results with simplified physics. *Quarterly Journal of the Royal Meteorological Society*, 126(564):1143–1170, 2000.
- Alain Robichaud. Surface data assimilation of chemical compounds over north america and its impact on air quality and air quality health index (aqhi) forecasts. *Air Quality, Atmosphere & Health*, 10:955–970, 2017.
- Yves J Rochon, Michael Sitwell, and Young-Min Cho. A study on harmonizing total ozone assimilation with multiple sensors. *Atmospheric Chemistry and Physics*, 19(14):9431–9451, 2019.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- Mark W Shephard, Enrico Dammers, Karen E Cady-Pereira, Shailesh K Kharol, Jesse Thompson, Yonatan Gainariu-Matz, Junhua Zhang, Chris A McLinden, Andrew Kovachik, Michael Moran, S. Bittman, C. E. Sioris, D. Griffin, M. J. Alvarado, C. Lonsdale, V. Savic-Jovcic, and Q. Zheng. Ammonia measurements from space with the crosstrack infrared sounder: characteristics and applications. *Atmos. Chem. Phys.*, 20(4):2277–2302, 2020. doi: https: //doi.org/10.5194/acp-20-2277-2020.

- MW Shephard and KE Cady-Pereira. Cross-track infrared sounder (cris) satellite observations of tropospheric ammonia. *Atmos. Meas. Tech.*, 8(3):1323–1336, 2015. doi: https://doi.org/10.5194/amt-8-1323-2015.
- MW Shephard, KE Cady-Pereira, M Luo, DK Henze, RW Pinder, JT Walker, CP Rinsland, JO Bash, L Zhu, VH Payne, and L Clarisse. Tes ammonia retrieval strategy and global observations of the spatial and seasonal variability of ammonia. *Atmos. Chem. Phys.*, 11(20):10743–10763, 2011. doi: https://doi.org/10.5194/acp-11-10743-2011.
- Michael Sitwell, Mark W Shephard, Yves Rochon, Karen Cady-Pereira, and Enrico Dammers. An ensemblevariational inversion system for the estimation of ammonia emissions using cris satellite ammonia retrievals. *Atmospheric Chemistry and Physics*, 22(10):6595–6624, 2022.
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. Advances in neural information processing systems, 33:12438–12448, 2020.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- Jean-Noéul Thepaut and Philippe Courtier. Four-dimensional variational data assimilation using the adjoint of a multilevel primitive-equation model. *Quarterly Journal of the Royal Meteorological Society*, 117(502):1225–1254, 1991.
- Anna Vaughan, Stratis Markou, Will Tebbutt, James Requeima, Wessel P. Bruinsma, Tom R. Andersson, Michael Herzog, Nicholas D. Lane, Matthew Chantry, J. Scott Hosking, and Richard E. Turner. Aardvark weather: end-toend data-driven weather forecasting. 2024. URL https://arxiv.org/abs/2404.00411.
- Cynthia H Whaley, Paul A Makar, Mark W Shephard, Leiming Zhang, Junhua Zhang, Qiong Zheng, Ayodeji Akingunola, Gregory R Wentworth, Jennifer G Murphy, Shailesh K Kharol, and K. E. Cady-Pereira. Contributions of natural and anthropogenic sources to ambient ammonia in the athabasca oil sands and north-western canada. *Atmos. Chem. Phys.*, 18(3):2011, 2018. doi: https://doi.org/10.5194/acp-18-2011-2018.
- Evan White, Mark W Shephard, Karen E Cady-Pereira, Shailesh K Kharol, Sean Ford, Enrico Dammers, Evan Chow, Nikolai Thiessen, David Tobin, Greg Quinn, Jason O'Brien, and Jesse Bash. Accounting for non-detects: Application to satellite ammonia observations. *Remote Sensing*, 15(10):2610, 2023.
- Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision* (ECCV), pages 3–19, 2018.
- Yi Xiao, Lei Bai, Wei Xue, Kang Chen, Tao Han, and Wanli Ouyang. Fengwu-4dvar: Coupling the data-driven weather forecasting model with 4d variational assimilation. *arXiv preprint arXiv:2312.12455*, 2023.
- L Zhu, DK Henze, KE Cady-Pereira, MW Shephard, M Luo, RW Pinder, JO Bash, and G-R Jeong. Constraining us ammonia emissions using tes remote sensing observations and the geos-chem adjoint model. *J. Geophys. Res.: Atmospheres*, 118(8):3355–3368, 2013. doi: https://doi.org/10.1002/jgrd.50166.
- Milija Zupanski. Regional four-dimensional variational data assimilation in a quasi-operational forecasting environment. *Monthly weather review*, 121(8):2396–2408, 1993.

# Supplement to EnsAI: An Emulator for Atmospheric Chemical Ensembles

# S1 Further Information on Training the Network

This section contains additional details of the training of the U-Net network. The network was trained using a batch size of 32 and learning rate of  $10^{-4}$  through 50 epochs. The learning curve for the training is plotting in Figure S1.



Figure S1: The learning curve produced during the training of the EnsAI model.

# S2 Additional Plots for the Direct Comparison of Covariances



Figure S2: Statistics for the standard deviation of surface ammonia  $\sigma_c$  taken over the combined evaluation time period of July 1-7 and August 1-7. The left panel shows the mean value of  $\sigma_c^{\text{GM}}$  over this time, while the middle and right panels show the  $RMSE(\sigma_c)$  values for the EnsAI ensemble and static covariance, respectively.



Figure S3: RMSE values for the horizontal correlation length modes  $L_{ec}[m]$  for  $m \leq 2$  for the emissions/surface concentration cross-correlation over the combined time period of July 1-7 and August 1-7 for the EnsAI ensemble (left column) and the static covariance (right column). A narrow band of grid points around the lateral boundaries of the model grid is excluded as the full mode expansion of  $L_{ec}$  cannot be computed at these points.