Master thesis proposal: Structured uncertainty estimation for video restoration

Supervisor: Pablo Arias Martínez (UPF)

January 22, 2026

This master thesis is about video restoration problems, such as noise removal and super-resolution. In most image restoration problems, we have a degraded observation y of an unknown ideal image x, and it is impossible to invert the degradation. There are in fact, infinite possible solutions which are compatible with the observed degradation y, as shown in Figure 1. These solutions are characterized by the posterior distribution.

In the last years deep neural networks have achieved remarkable advances in restoration quality. A problem of this type of methods is that -for the most part- we do not understand how they work: they are considered black box approaches which estimate a single solution, and for which it is difficult to provide guarantees on their performance or to characterize their failure modes. These are important drawbacks in safety-critical applications.

One way to address this issue is by estimating the uncertainty of the result. There are different ways to quantify uncertainty. We will consider the variance of the distribution of possible solutions. This tells us how far away the unknown real solution could be from the output given by the network (i.e. the size of the expected error).

Using appropriate loss functions, we can train an image restoration network to predict the *per-pixel* variance of the posterior distribution (an example is shown in Figure 2). But these per-pixel variances fail to capture important correlations between pixels.

Recently, some methods have proposed to estimate the covariance matrix of the posterior distribution[1] This provides the principal components of the space of possible solutions, which can be used to navigate the distribution of possible solutions by drawing approximate samples from the posterior distribution.

[1] *Dorta, G., et al. Structured uncertainty prediction networks. CVPR 2018.*

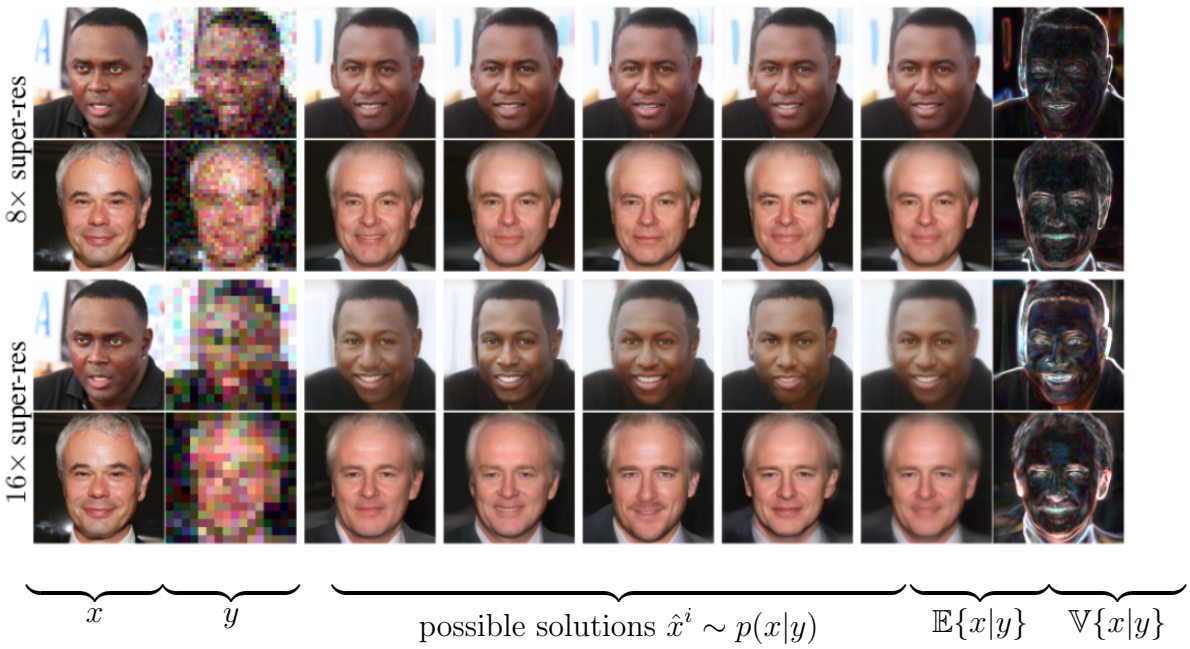

Figure 1: *An image restoration problem:* Given a degraded observation y there are many possible solutions carachterized by the posterior distribution $p(x|y)$.

This interesting not only for visualizing possible solutions, but also as an important source of information for other tasks that are computed based on the restored image (such as a segmentation algorithm, a detection algoritm, etc).

The goals of this thesis are the following:

Primary Goal: Train image restoration networks to predict **structured uncertainty** to better capture the distribution of possible solutions.

Secondary Goal: Explore the use of this structured uncertainty on a downstream task (segmentation, detection, restoration of a future frame). This last point would be novel and could lead to a publication.

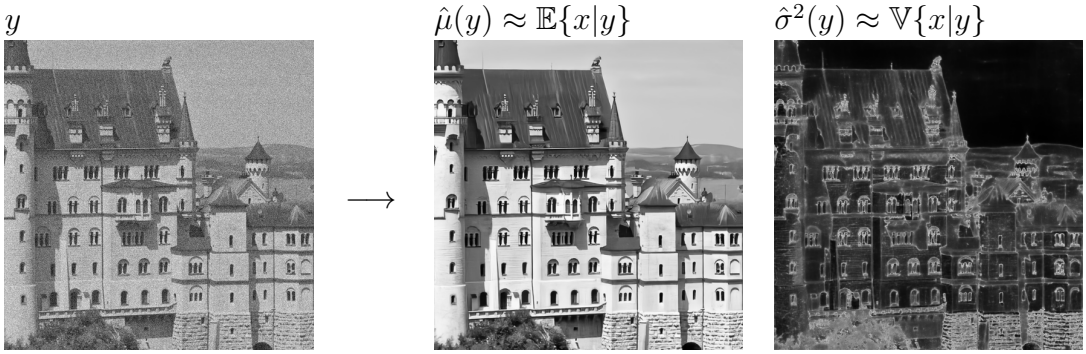

Figure 2: The denoised image and with the per-pixel variance map produced by a network trained with uncertainty quantification.

.