

Dust in a compact, cold, high-velocity cloud: A new approach to removing foreground emission

D. Lenz, L. Flöer, and J. Kerp

Argelander-Institut für Astronomie (AIfA), Universität Bonn, Auf dem Hügel 71, 53121 Bonn, Germany
e-mail: dlenz@astro.uni-bonn.de

Received 13 April 2015 / Accepted 13 December 2015

ABSTRACT

Context. Because isolated high-velocity clouds (HVCs) are found at great distances from the Galactic radiation field and because they have subsolar metallicities, there have been no detections of dust in these structures. A key problem in this search is the removal of foreground dust emission.

Aims. Using the Effelsberg-Bonn HI Survey and the *Planck* far-infrared data, we investigate a bright, cold, and clumpy HVC. This cloud apparently undergoes an interaction with the ambient medium and thus has great potential to form dust.

Methods. To remove the local foreground dust emission we used a regularised, generalised linear model and we show the advantages of this approach with respect to other methods. To estimate the dust emissivity of the HVC, we set up a simple Bayesian model with mildly informative priors to perform the line fit instead of an ordinary linear least-squares approach.

Results. We find that the foreground can be modelled accurately and robustly with our approach and is limited mostly by the cosmic infrared background. Despite this improvement, we did not detect any significant dust emission from this promising HVC. The 3σ -equivalent upper limit to the dust emissivity is an order of magnitude below the typical values for the Galactic interstellar medium.

Key words. ISM: clouds – dust, extinction – methods: data analysis

1. Introduction

Since their discovery by Muller et al. (1963), high-velocity clouds (HVCs) have been the target of a wide range of studies (see Wakker & van Woerden (1997) for a review). It is thought that they are located at distances of several kpc Wakker 2001 and contribute to the fuelling of low-metallicity gas into the Galaxy (Putman et al. 2012).

Early attempts to detect dust emission from HVCs were unsuccessful (Wakker & Boulanger 1986; Desert et al. 1988) and were generally considered to be difficult because of the clouds' low metallicities and hence low dust-to-gas ratios (Fox et al. 2004). Moreover, HVCs are located far from the interstellar radiation field (ISRF), which implies a faint illumination by UV light that is absorbed and re-emitted by the dust grains. The cosmic infrared background radiation (CIB) has anisotropies on angular scales that are comparable to the typical sizes of HVCs and is therefore another source of confusion (Planck Collaboration XXIV 2011; Planck Collaboration XXX 2014). Very recently, the Planck Collaboration Int. XVII (2014) has reported that the variation of dust emissivities across the field of interest is the limiting source of uncertainty when modelling the dust data.

Recent attempts to disclose the far-infrared (FIR) emissivity of HVCs are in line with these findings. Neither the investigation of different high-latitude clouds (Planck Collaboration XXIV 2011) nor the stacking of GALFA-HI compact clouds (Saul et al. 2014) has detected significant FIR emission.

Despite these odds, Miville-Deschênes et al. (2005b) generate a simple model of the dust emission from HVC complex C based on its HI column density and find a faint, but significant dust emissivity for the HVC regime. A similar approach,

combined with an investigation of the chance correlation of HI and dust emission, yields a $>3\sigma$ detection of dust towards complex M (Peek et al. 2009).

Here we investigate the HVC located at $(l, b, v_{\text{LSR}}) = (125^\circ, 41^\circ, -207 \text{ km s}^{-1})$, hereafter HVC125. The cloud has been previously studied with the Effelsberg 100 m telescope (Brüns et al. 2001) and the Westerbork Synthesis Radio Telescope (Braun & Burton 2000). In combination, these data sets disclose a two-phase structure and a head-tail morphology of the HVC of interest. This might indicate ram-pressure interaction which in turn results in a reduced formation time of H_2 and dust (Guillard et al. 2009; Röhser et al. 2014). The single-dish data suggest that the warm phase is stripped off the core of the HVC. The cold component has an extraordinarily narrow line width around $\Delta v = 2 \text{ km s}^{-1}$ FWHM, equivalent to a kinetic temperature of $T_{\text{kin}} \lesssim 85 \text{ K}$. The compact spatial structure of a few arcminutes, high brightness temperature of $T_{\text{B}} \gtrsim 10 \text{ K}$ in the single-dish data, and low kinetic temperature make HVC125 one of the most promising HVCs in terms of detection probability of FIR dust emission.

2. Data

We used data from the recently finished Effelsberg-Bonn HI Survey (EBHIS, Kerp et al. 2011; Winkel et al. 2010, 2016) to study the neutral atomic hydrogen in HVC125.

For the FIR data, we used the latest release of the *Planck* data at 857 GHz (Planck Collaboration VIII 2016). We chose the highest frequency of the *Planck* data because the relative contributions of the cosmic microwave background and the cosmic infrared background to the uncertainty in modelling the fore-

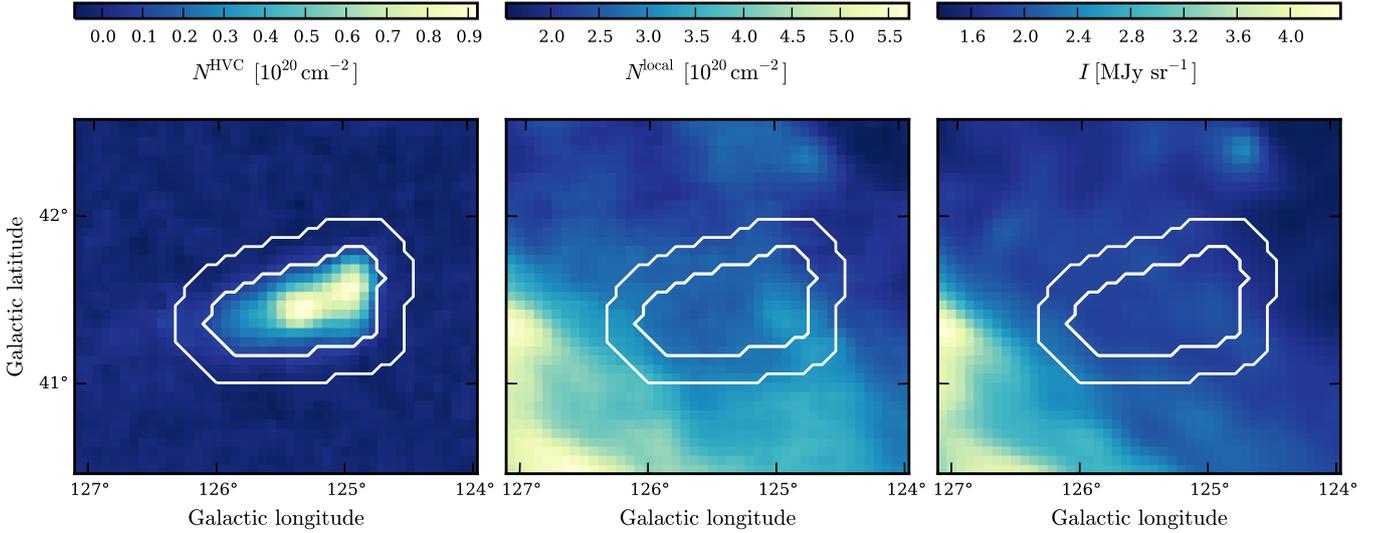


Fig. 1. *Left:* EBHIS column density map N^{HVC} of the HVC ($-230 \text{ km s}^{-1} < v_{\text{LSR}} < -190 \text{ km s}^{-1}$). *Center:* EBHIS column density map N^{local} of the local foreground emission ($-190 \text{ km s}^{-1} < v_{\text{LSR}} < +30 \text{ km s}^{-1}$). *Right:* FIR intensity $I_{857 \text{ GHz}}$ from [Planck Collaboration VIII \(2016\)](#). The inner white contour line corresponds to the 5σ noise level in the HVC column density map. The outer contour marks an annulus that contains the same number of pixels as the tight HVC mask.

ground both decrease in proportion to the frequency ([Planck Collaboration Int. XVII 2014](#), their Fig. C.1).

To account for the differences in angular resolution, the dust data were smoothed to the angular resolution of the EBHIS data by Gaussian convolution. The final angular resolution is $10.83'$. The EBHIS data have a spectral resolution of 1.49 km s^{-1} at a channel spacing of 1.28 km s^{-1} .

3. Analysis

In the following, we use N to refer to the HI column density N_{HI} and I to refer to the FIR intensity at 857 GHz. When analysing the dust content of HVCs by comparing their HI column density to their dust content, the most challenging step is the estimation of the foreground dust emission. Because of the complexity and uncertainty of the correlation of dust and gas, an accurate and robust determination of the foreground component is of the utmost importance (e.g. [Peek et al. 2009](#); [Saul et al. 2014](#)).

The HI data allow us to distinguish between local foreground emission and HVC emission via the radial velocity. For HVC125, we selected the velocity range ($-230 \text{ km s}^{-1}, -190 \text{ km s}^{-1}$) for the HVC and the remaining range ($-190 \text{ km s}^{-1}, +30 \text{ km s}^{-1}$) for the foreground emission. The corresponding column density maps are shown in Fig. 1. Moreover, we present an image of the FIR intensity at 857 GHz in the direction of HVC125. The inner contour outlines the 5σ level of the HVC HI column density. The outer contour marks an annulus around the HVC that contains as many pixels as the inner, narrow mask. In the following, we use these masks to determine the foreground dust contribution and the dust emissivity of the HVC.

For this analysis, we decided not to include the statistical uncertainties from the data noise of $\sigma_{\text{rms}} = 90 \text{ mK}$ for the EBHIS data and $\sigma_{\text{rms}} = 0.014 \text{ MJy sr}^{-1}$ for the *Planck* data at 857 GHz. We refer again to Fig. C.1 in [Planck Collaboration Int. XVII \(2014\)](#), which shows that at this frequency the analysis is dominated by uncertainties from the foreground estimation.

Table 1. Emissivities ϵ^i for the different velocity components and offset Z , according to Eq. (1).

Mask	ϵ^{local}	ϵ^{HVC}	Z
Small	0.26 ± 0.48	0.05 ± 0.38	1.14 ± 1.31
Wide	0.41 ± 0.09	0.029 ± 0.11	0.76 ± 0.26
None	0.61 ± 0.005	-0.05 ± 0.03	0.26 ± 0.02

Notes. Units are $\text{MJy sr}^{-1}/10^{20} \text{ cm}^{-2}$ for the emissivities and MJy sr^{-1} for the offset. The fit uncertainties are asymptotic standard errors, taken from the covariance matrix.

3.1. Standard approach

The standard approach (e.g. [Miville-Deschênes et al. 2005b](#); [Planck Collaboration XXIV 2011](#)) used to evaluate the dust content of HVCs is the superposition of different HI column density maps N^i to model the FIR intensity:

$$I(x, y) = \epsilon^{\text{local}} N^{\text{local}}(x, y) + \epsilon^{\text{HVC}} N^{\text{HVC}}(x, y) + Z. \quad (1)$$

Here, I is the FIR intensity, ϵ denotes the dust emissivity per hydrogen nucleon, and Z is a constant offset; I and N^i are two-dimensional images with (x, y) denoting the spatial position of the pixel.

We fitted Eq. (1) with a least-squares approach to quantify the parameters ϵ^{local} , ϵ^{HVC} , and Z . To investigate the influence of the spatial area that is fitted on the results, we performed the fit for three different spatial masks: a tight mask around the HVC, a more extended mask, and the full image. The results are compiled in Table 1. The large variations for different spatial masks and the uncertainties on the fit parameters emphasise that this approach is very sensitive to changes in the area of interest. Furthermore, this approach only relies on two parameters (Eq. (1)) to describe the foreground: the local dust emissivity ϵ^{local} and the offset Z and so it cannot account for multiple features at different radial velocities, possibly exposed to different physical environments. For the full field, we show the modelled FIR intensity I^{Standard} and the residual $I - I^{\text{Standard}}$ in Fig. 3, left column.

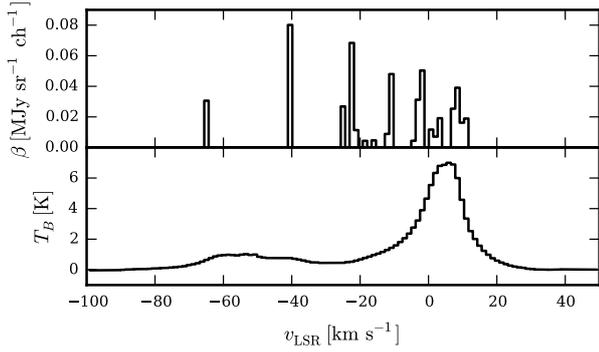


Fig. 2. *Top:* GLM coefficients β^i for each channel based on the cross-validated lasso regression (Eqs. (2) and (3)). *Bottom:* mean HI spectrum of the data cube.

3.2. Generalised linear model for foreground estimation

To overcome the limitations of the standard approach, we applied a generalised linear model (GLM; for a review see Madsen & Thyregod 2010; de Souza et al. 2014) to the data. For this, we assume that each channel T_B^i of the HI cube can contribute individually to the FIR intensity. In consequence, we do not rely on a vague definition of the velocity range for local and HVC gas.

Within the GLM, the FIR intensity can be written as

$$I^{\text{local}}(x, y) = \sum_i T_B^i(x, y)\beta^i + Z. \quad (2)$$

The β^i are the GLM coefficients and can be understood as emissivity per spectral channel. The parameter Z is a global offset to the model. Because of the co-linearity between neighbouring spectral channels, the assumption of independent data for ordinary least-squares fitting is violated. To break this degeneracy, we controlled the GLM with lasso regularisation (Tibshirani 1996) and minimised the term

$$\|I^{\text{local}} - T_B^i\beta^i - Z\|^2 + \alpha \cdot \|\beta^i\|_1. \quad (3)$$

The first part of this term corresponds to the regular least-squares approach and gives the residual sum of squares. The second part is the penalty term and ensures that the coefficients β^i are chosen as sparsely as possible. The strength of this second term is scaled by α . We used simulations and cross validation to optimise this regularisation strength (see Sect. 4). The regularised GLM and the cross validation was implemented via `scikit-learn` (Pedregosa et al. 2011), a machine-learning package for python.

For our application to the HI and dust data, only data points outside of the wide mask and with $v_{\text{LSR}} > -190 \text{ km s}^{-1}$ were considered because they are unrelated to the HVC. This ensures that the HVC signal is not accidentally removed by our foreground subtraction. While the distinction between HVC and local gas is very straightforward in this particular case, there are other cases where gas at high or intermediate velocities is difficult to disentangle (e.g. HVC Complex M, Wakker 2001). In these cases, we cannot simply remove the foreground by applying a threshold in radial velocity, but have to rely only on spatial masking of the HVC.

The resulting GLM coefficients β^i and the mean HI spectrum are shown in Fig. 2. The offset is $Z = 0.73 \text{ MJy sr}^{-1}$. We find that the dust emission towards HVC125 can be modelled well with approximately seven different emissivities. The narrowness of the GLM coefficients is the product of the applied regularisation. Visually inspecting the HI data cube discloses that

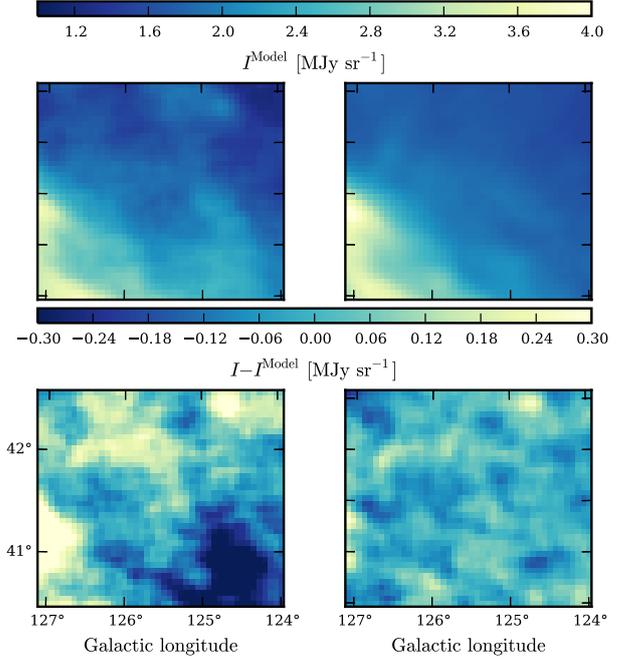


Fig. 3. Model of the FIR intensity (*top*) and residual emission (*bottom*). We show the results of the standard approach (Sect. 3.1) in the *left column* and the GLM (Sect. 3.2) in the *right column*.

each individual GLM coefficient is indeed associated physically with an individual HI structure. However, the exact matching of HI structures to FIR emissivity peaks varies with the regularisation strength and demands further justification (Sect. 4). The sparse filling of the whole HI spectrum with GLM coefficients argues strongly for a localisation of the dust within the cold neutral medium (CNM) filaments and is inconsistent with a homogenous mixture of dust and gas on all linear scales. However, we note that the discrete nature of the FIR emissivities for each spectral channel is primarily a schematic description.

Similar to the standard approach, we show the modelled FIR intensity and the residual (Fig. 3, right column). A comparison of the two different approaches shows that the GLM can cover a wider range of FIR intensities and manages to account for a larger number of different features. The residual is less structured and the remaining structures stem from the CIB anisotropies (compare Sect. 4).

Furthermore, we show the histograms from the residual FIR intensity in Fig. 4. For both models, we consider only data points outside of the wide mask because we are only interested in the capability to model the foreground emission. For the standard approach, we find that the residual is rather broad and irregular. In contrast, the GLM residual is narrow and symmetric and its shape can be approximated as Gaussian.

3.3. Measurement of the HVC dust content

To quantify the hypothetical dust emission from HVC125, we corrected the FIR intensity map for the local foreground emission by using the GLM, $I^{\text{HVC}} = I - I^{\text{GLM}}$ (Fig. 3, bottom right). We investigated the correlation between the HVC HI column density and the HVC FIR intensity within the narrow mask with a linear correlation in the Bayesian framework (e.g. D'Agostini 2003). Thus, we sampled the posterior given by

$$p(\epsilon, \sigma | \mathcal{D}) \propto \mathcal{L}(\mathcal{D} | \epsilon, \sigma) p(\epsilon, \sigma) \quad (4)$$

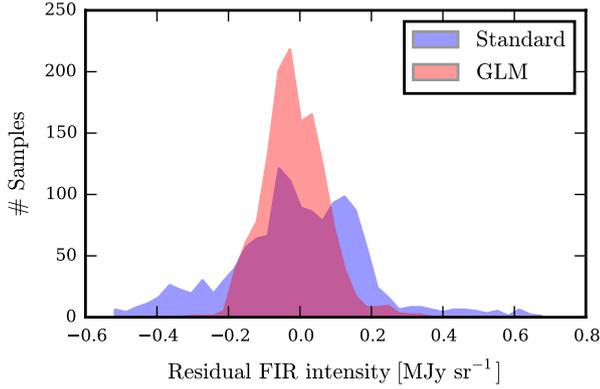


Fig. 4. Residual histograms after correcting for the local foreground FIR intensity outside of the wide mask for the two different methods.

where \mathcal{D} is the data vector $\mathcal{D} = (N^{\text{HVC}}, I^{\text{HVC}})$. The likelihood is given by

$$\mathcal{D}|\epsilon, \sigma \sim \mathcal{N}(I^{\text{HVC}} - \epsilon \cdot N^{\text{HVC}}, \sigma^2). \quad (5)$$

The parameter ϵ is the dust emissivity per hydrogen nucleus and σ is the standard deviation of the Gaussian likelihood.

We selected minimally informative priors for our parameters:

$$\begin{aligned} \epsilon &= \tan \phi \\ p(\phi) &= \text{Uniform}(\phi, 0, \pi/2) \\ p(\log \sigma) &= \text{Uniform}(\log \sigma, -8, 4). \end{aligned} \quad (6)$$

For the emissivity ϵ , we sampled uniformly in $\arctan(\epsilon)$ to avoid a bias towards larger values. A simple scale-invariant prior was chosen for the scatter σ (Jeffreys 1946) in a reasonable range. Furthermore, we reasonably chose to restrict ϵ to positive values. An offset FIR intensity is already part of the GLM, hence we choose not to include a further offset here. Lastly, we note that we did not account for the spatial covariance in the image owing to the non-flat CIB angular power spectrum. For a complete and proper treatment of this effect, we would also require an accurate determination of the spatial covariance due to the beam shapes of the different data sets and the sampling on the pixel grid.

The model was sampled with `emcee` (Foreman-Mackey et al. 2013), a python implementation of the affine-invariant ensemble sampler for Markov chain Monte Carlo (MCMC, Goodman & Weare 2010). Figure 5 shows randomly drawn samples of this model, applied to the narrow mask HI and dust data of HVC125. The posterior distribution of the individual parameters is presented in Fig. 6.

As in the standard approach (Table 1), we not only evaluated the narrow mask, but also the wide mask and the full field. The resulting fit parameters are summarised in Table 2.

We find that the emissivity ϵ and scatter σ parameters are well-sampled. The HVC emissivity is not normally distributed and illustrates that the model strongly prefers zero emissivity. The 99.87% upper limit, corresponding to 3σ , is $0.021 \text{ MJy sr}^{-1}/10^{20} \text{ cm}^{-2}$ and thus an order of magnitude below typical Galactic ISM values (Planck Collaboration XXIV 2011). Moreover, the parameters vary only slightly for different masks. As expected, the scatter increases with the size of the mask, or equivalently with the number of data points.

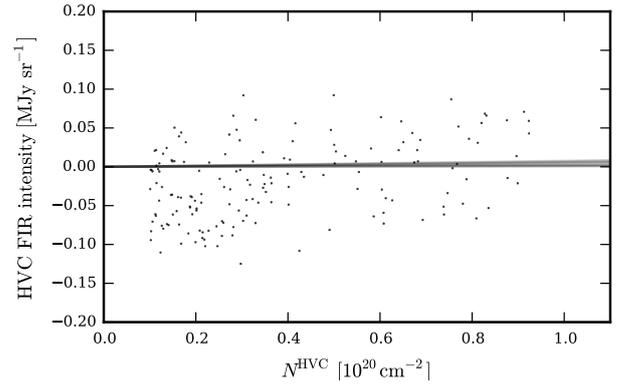


Fig. 5. Linear correlation between HVC HI column density and foreground-subtracted FIR intensity in the narrow mask. The lines correspond to thirty randomly chosen MC samples (Eqs. (4) to (6)) from the Bayesian line fit.

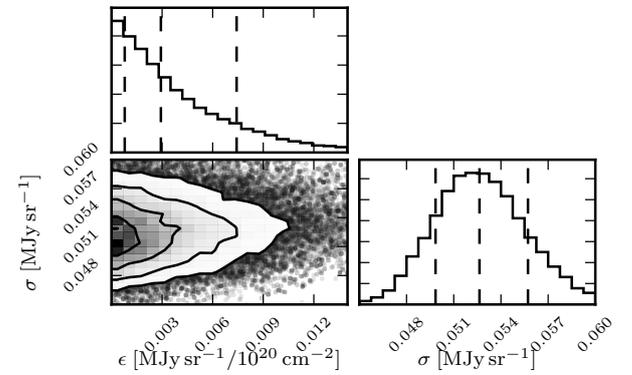


Fig. 6. Posterior distribution for ϵ and σ , obtained from sampling Eq. (4). The lines in the histogram indicate the 16th, 50th, and 84th percentile, equivalent to mean and $\pm 1\sigma$ for a Gaussian posterior.

Table 2. Emissivity ϵ and scatter σ for the different masks.

Mask	ϵ [$\text{MJy sr}^{-1}/10^{20} \text{ cm}^{-2}$]			σ [MJy sr^{-1}]		
	16%	Median	84%	16%	Median	84%
Small	0.001	0.003	0.007	0.049	0.053	0.056
Wide	0.001	0.002	0.006	0.045	0.047	0.049
None	0.001	0.003	0.008	0.064	0.065	0.066

Notes. The numbers indicate the median of the posterior distribution and the 16th and 84th percentile. For a Gaussian posterior, this is equivalent to median $\pm 1\sigma$.

4. Verification of the GLM

In the following, we present our thorough investigation of the performance of the GLM; we conducted a series of simulations and tests to explore its advantages and limitations. Moreover, we used these simulations to properly select the regularisation strength α (Eq. (3)).

4.1. Construction of the simulations

The simulations were generated in the following way:

1. We generated an artificial, noisy spectrum of GLM coefficients;
2. we convolved this coefficient spectrum with the measured HI data cube to generate a map of foreground FIR intensity;
3. We added a random realisation of the CIB to the foreground.

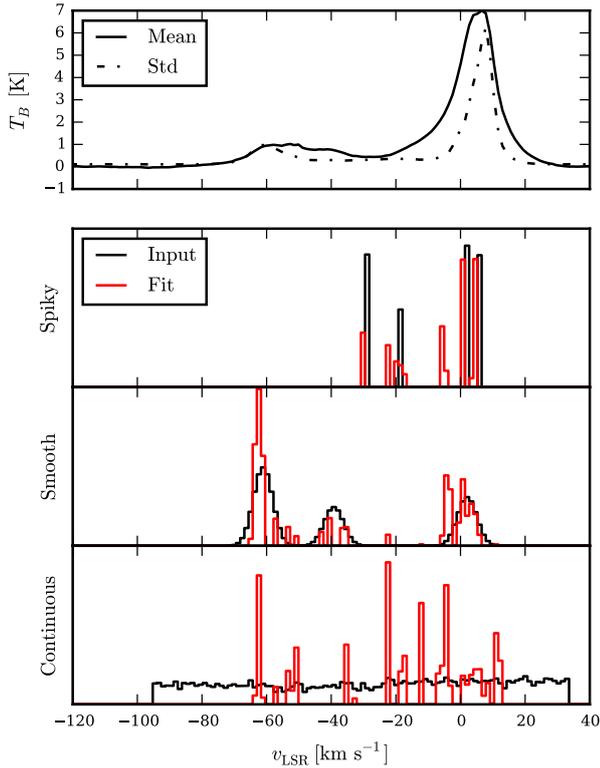


Fig. 7. *Top panel:* spectrum of mean and standard deviation of the full HI data cube. *Bottom panels:* input spectra (black) of GLM coefficients to simulate FIR intensity maps. The reconstruction by the GLM is shown in red and has been shifted to the left by one channel for illustration purposes. See the text for a detailed description.

Because we did not know the real nature of dust emissivities at different radial velocities, we tested different approaches. For all of them, we assumed that the dust emissivity occurs only in spectral channels in which significant HI emission is found. We distinguished between the following types of GLM coefficients (see also Fig. 7):

Spiky: one spectral channel wide, between 2 and 7 components. The spikes vary in amplitude by up to 40%.

Smooth: similar to the spiky input, but the spectrum of GLM coefficients is smoothed with random Gaussian kernels with FWHM between 2 km s^{-1} and 8 km s^{-1} .

Flat: this represents the standard approach to investigating the HI-FIR correlation. We choose a constant dust emissivity with 10% random fluctuation for local gas ($v_{\text{LVC}} > -30 \text{ km s}^{-1}$) and a 20% lower dust emissivity for the gas at intermediate velocities. ($-95 \text{ km s}^{-1} < v_{\text{LVC}} < -30 \text{ km s}^{-1}$)

The amplitudes of all these GLM coefficient spectra were normalised such that the resulting mean intensity of the foreground FIR map was equal to the value found in the measured FIR intensity map (Fig. 1, right panel). This ensured that we modelled the proper ratio of foreground to background.

To combine the foreground FIR intensity with the CIB, we generated random realisations of the Gaussian random field based on the CIB angular power spectrum taken from Planck Collaboration XXX (2014). We extrapolated the angular power spectrum from their Table D.2 with a power-law. This extrapolation does not hold for large angular separations i.e. small multipoles, but this effect is negligible for the present field size of only $2^\circ \times 3^\circ$. Here, a power law is a valid approximation

(G. Lagache, priv. comm.). To obtain the proper ratio of foreground and background components, we also scaled the simulated CIB to match the mean and fluctuation amplitude given in Planck Collaboration XVIII (2011, their Table 5). Despite the non-flat angular power spectrum of the CIB, this is possible because the angular size of the fields investigated in Planck Collaboration XVIII (2011) is similar to the field size in the present study. Finally, the simulated CIB was smoothed to the angular resolution of $10.8'$, which was used for all data throughout this study.

We generated 1000 of these simulations for each type of spectrum (spiky, smooth and flat) and applied the GLM to reconstruct the total FIR intensity. We compare the input and outcome of the GLM coefficient spectra in Fig. 7.

4.2. GLM performance on simulations and choice of regularisation strength

The visual inspection of Fig. 7 shows that, despite the CIB confusion, we were able to properly reconstruct the shape and position of the GLM coefficients for the spiky and smooth case. Because of the GLM design, we did not recover the exact shape of the flat input spectrum, but used the most relevant channels to create an approximation that produces an accurate model. We note again that the discrete description of the dust emissivities for the spectral channels is a result of our approach and does not necessarily reflect the physical conditions.

The results are furthermore evaluated via three quantities: the mean strength of the CIB, the strength of its fluctuations and the Pearson's r correlation coefficient of the reconstructed and the input CIB image. This mainly ensures that the GLM neither over- nor underfits the data. For the CIB mean μ and fluctuation amplitude σ , the values are taken from Planck Collaboration XVIII (2011).

We investigated how these quantities vary as a function of the regularisation strength α (Fig. 8). Our example was generated for the smooth input GLM coefficients. For the other cases, the results are presented in Sect. A. Dashed lines indicate the input values, solid lines are the results of our application of the GLM to the simulated data. The contours correspond to 1σ uncertainties.

We find that for $\alpha \lesssim 4 \times 10^{-3}$, there is an agreement between input and reconstruction for all three estimators within their respective uncertainties. We chose to set $\alpha = 2 \times 10^{-3}$ for our analysis of HVC125 (vertical line in Fig. 8).

The Pearson's r correlation coefficient between the input and reconstructed CIB is remarkably constant and close to 1, even for a very weak regularisation that allows a great number of GLM coefficients. This illustrates that chance correlation by individual HI spectral channels is not very efficient in mimicking the CIB signal. For a very strong regularisation, the model cannot account for all the dust-emitting HI components and the residual map is dominated by foreground emission, yielding a poor correlation to the input CIB.

The CIB mean μ is the quantity that varies strongest for different values of the regularisation strength α . We find that for $\alpha = 10^{-3}$ to $\alpha = 4 \times 10^{-3}$, the CIB mean is properly estimated.

The CIB fluctuation amplitude σ is systematically underestimated if the regularisation is too weak, meaning that the model overfits the data. Other than in the Pearson's r correlation, this hints towards a mimicking of background CIB by chance correlation with some HI channels. This effect, however, is of the order of 10% and is within the uncertainties for our choice of $\alpha = 2 \times 10^{-3}$.

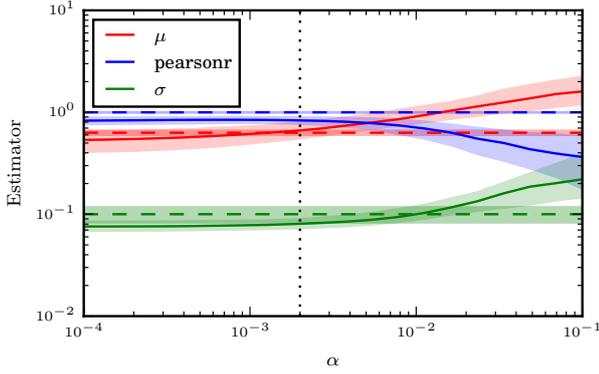


Fig. 8. Evaluation of the reconstructed CIB mean μ , CIB fluctuation amplitude σ and Pearson's r of input CIB image and reconstruction. This is based on the smooth input GLM coefficients (third panel from the top in Fig. 7). Dashed lines indicate the input, solid lines the reconstructed quantities. Contours correspond to 1σ uncertainties. The vertical line indicates our choice of α .

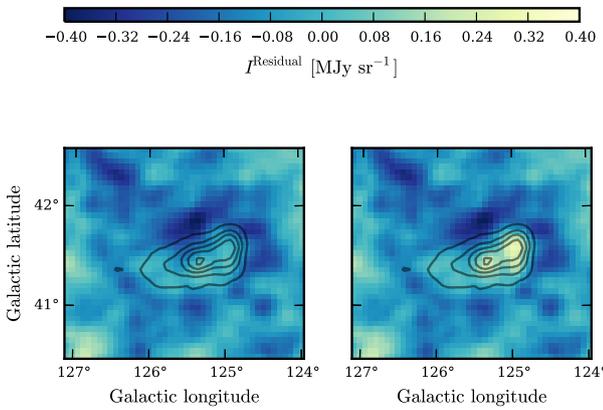


Fig. 9. GLM reconstruction of a simulated dusty HVC. The image is dominated by the simulated CIB anisotropies. The HVC dust emissivity corresponds to 10% (left) and 30% (right) of typical Galactic values. The black contours correspond to the HVC column density starting at 10^{19} cm^{-2} and increasing in steps of $2 \times 10^{19} \text{ cm}^{-2}$.

Furthermore, we find that the choice of α and the evaluation of the different quality estimators does not vary strongly for different types of input GLM coefficients (Figs. A.1, A.2). Because we cover a variety of shapes (Fig. 7) and demonstrate that the GLM can properly remove the foreground and uncover the faint CIB emission in each individual case, we conclude that it is well suited for the search of dust in HVCs.

To conclude, our approach is not constructed to precisely measure the dust emissivity of the individual clouds and filaments; rather, it is designed to remove the local foreground for studies of faint FIR signals such as CIB emission or dust in HVCs. To further verify this, we simulated a FIR intensity map in which the HVC has a dust emissivity of 10% and 30% of typical Galactic values and removed the foreground emission (Fig. 9). We find that for an emissivity of only 10% Galactic, the HVC can barely be distinguished from the residual CIB fluctuations. For the higher emissivity, the HVC signal is significantly stronger than these fluctuations and would imply a detection in our Bayesian line fit.

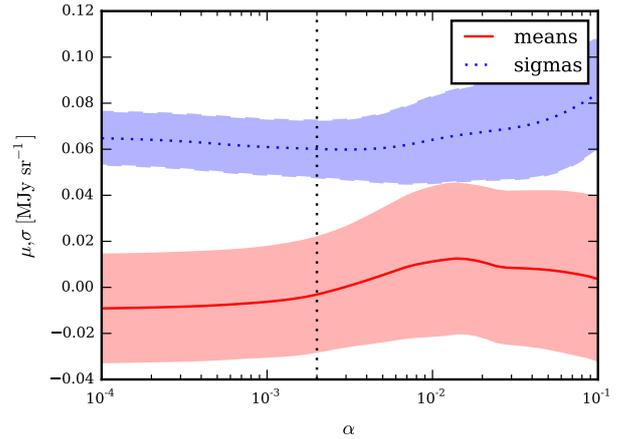


Fig. 10. Results of the cross validation to investigate our choice of the regularisation strength. The solid lines correspond to the mean and standard deviation of the residual emission in the test sample. Contours correspond to 1σ uncertainties. The vertical line illustrates our choice of α .

4.3. Estimation of alpha and the uncertainties via cross validation

To further validate the regularisation strength α in Eq. (3), we used structured cross validation (e.g. Picard & Cook 1984) with the HVC mask as kernel. For this purpose, we shifted the wide HVC mask to random positions in our field. The GLM was then computed on the data outside of this mask (training sample). Subsequently, we compared the data and the model inside the mask (test sample) by inspecting the residual mean and its standard deviation. This was done for 1000 different, random HVC mask positions. The resulting means and 1σ uncertainties of residual mean and residual standard deviation in the test sample are shown in Fig. 10 for a range of regularisation strengths.

If the regularisation is too strong, the model will eventually fail and will not properly account for the complexity of the data. Based on the simulations, we chose $\alpha = 2 \times 10^{-3}$. This is supported by the findings of the present cross validation which shows that this value is the strongest possible regularisation before the model begins to lose accuracy.

We used the same technique to estimate the uncertainties for a fixed value of the regularisation strength $\alpha = 2 \times 10^{-3}$ and find that the mean residual FIR intensity is $-0.003 \pm 0.025 \text{ MJy sr}^{-1}$. The standard deviation, equivalent to the amplitude of the CIB anisotropies for the present field-size is $0.060 \pm 0.012 \text{ MJy sr}^{-1}$.

The same cross validation analysis for the standard approach yields a mean of $-0.040 \pm 0.129 \text{ MJy sr}^{-1}$ and a standard deviation of $0.129 \pm 0.054 \text{ MJy sr}^{-1}$. The larger uncertainties with respect to the GLM illustrate another advantage of the novel method. We note, however, that the standard deviation of the residuals differs significantly between the two methods. We discuss the implications further in Sect. 5.1.

4.4. HVC dust emissivities in simulations

To cross-check our result on the HVC dust emissivity (Sect. 3.3), we quantified its posterior distribution using simulated data in which no HVC dust emission is present. To generate the simulated FIR intensity maps, we used a smooth shape for the GLM coefficients (Fig. 7). After removing the foreground emission, we evaluated the HVC dust emissivity using a similar procedure to the one described in Sect. 3.3. We present the stacked posterior distribution for all 1000 simulations in Fig. 11. The 16th,

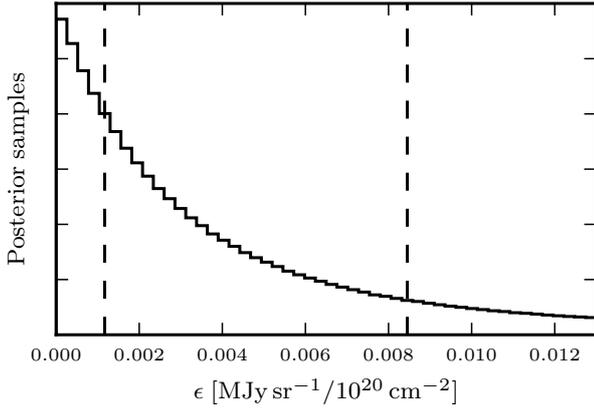


Fig. 11. Stacked posterior distribution for the HVC dust emissivity ϵ from simulated data. The procedure is similar to the one described in Sect. 3.3. The lines in the histogram indicate the 16th and 50th percentile.

50th, 84th and 99.87th percentiles correspond to a dust emissivity of 0.001, 0.008, 0.111, and 0.555 $\text{MJy sr}^{-1}/10^{20} \text{ cm}^{-2}$, respectively.

Despite the absence of a HVC signal in these simulations, the emissivities are lower than the values derived from real data (Fig. 6). The large values for the 84th and 99.87th percentile are the consequence of the large sample of simulations and yield a heavy tail in the posterior.

5. Discussion

5.1. Quality of the foreground model

To address the potential dust content of HVC125, a reliable estimation of the foreground FIR intensity is the most challenging step and of the utmost importance.

After early studies of the HI-dust correlation did not account for this foreground emission (Wakker & Boulanger 1986), the often-applied standard approach assumes a simple linear relation between FIR intensity and HI column density for different HI column density maps (Eq. (1), e.g. Miville-Deschênes et al. 2005b; Planck Collaboration XXIV 2011). This approach is limited by the uncertain separation of HI in different foreground and cloud components. Furthermore, it only has a low number of degrees of freedom and hence cannot properly treat the multiple, complex features at different radial velocities. Thus, clouds with different emissivities or illumination by the ISRF cannot be accounted for if they are in the same column density map. In the case of HVC125, the separation between HVC emission and foreground emission is straightforward and the results are not affected by uncertainties introduced by different definitions of N^{HVC} and N^{local} . Commonly, a clean separation of the different components for the more complex intermediate-velocity and local gas is hardly feasible. The residual histogram (Fig. 4) shows that the standard approach yields a broad and asymmetric residual signal. Moreover, Table 1 shows that the fit parameters vary strongly and are poorly determined for different spatial masks.

We apply a GLM to use each individual channel of the HI observations, restricted by the regularisation to overcome the degeneracy between the individual, correlated channels and to avoid overfitting. We justified the choice of the regularisation strength via simulations and cross validation, yielding consistent results. The success is well demonstrated by the narrow, symmetric distribution in the FIR intensity histogram. The results

of our cross validation yields no bias with a mean residual of $0.003 \pm 0.025 \text{ MJy sr}^{-1}$. The mean standard deviation of the residual maps of $0.060 \pm 0.012 \text{ MJy sr}^{-1}$ result from the underlying CIB fluctuations (Planck Collaboration XVIII 2011; Planck Collaboration XXX 2014). Moreover, the parameters are barely affected by different mask sizes, unlike the standard approach (Table 2). The spectrum of GLM coefficients illustrates that the local HI does not contribute uniformly to the foreground FIR intensity, but different features need to be accounted for. Aside from the complex, bright emission from the gas around 0 km s^{-1} , we find different filaments at intermediate velocities. Within the GLM framework, these features are recognised and the FIR intensity is successfully modelled.

The investigation of the standard deviation via cross validation generates significantly different results between the standard approach and the GLM. Here, the former is in agreement with studies of high-latitude fields of similar sizes using the very same approach (Planck Collaboration XVIII 2011; Planck Collaboration XXX 2014). However, our simulations show that we can properly reconstruct the different simulated CIB properties with the GLM approach. Accordingly, the tension in the strength of the CIB fluctuations requires further investigation in a future study. We note at this point that we constrain ourselves to this particular field on the sky which can only provide very limited insights into the global CIB properties.

5.2. Dust content of the HVC

We find that the foreground-corrected FIR intensity map does not contain any significant contribution from the HVC. This is clearly seen in the posterior distribution of the emissivity ϵ (Fig. 6, top), which strongly favours zero dust emissivity. Given our model, there is a 99.87% probability that the emissivity of the HVC is below $0.02 \text{ MJy sr}^{-1}/10^{20} \text{ cm}^{-2}$. This is an order or magnitude lower than typical Galactic values found by Planck Collaboration XXIV (2011).

Similar non-detections have been obtained in other studies of dust in HVCs (Planck Collaboration XXIV 2011; Saul et al. 2014). A noteworthy exception is HVC complex M (Peek et al. 2009). However, complex M is on the transition of HVC/IVC classification and could possibly be related to the IV arch (Wakker 2001).

Because of the very low kinetic temperature, its compact structure, and high brightness temperature, HVC125 is one of the most promising candidates for the detection of dust in HVCs. Moreover, the head-tail structure can indicate the formation of dust and molecules via the increased pressure (Gillmon et al. 2006; Guillard et al. 2009; Röhser et al. 2014). Our non-detection shows that even for this candidate, the upper limit is significantly below typical ISM dust emissivities.

6. Conclusion

6.1. Summary

To explore the properties and the origin of HVCs, their potential dust content can be a powerful tool. We pointed out the importance and the difficulty of estimating the foreground contribution to the FIR intensity in the classical framework. Without an accurate and robust determination of this foreground emission, it is not possible to evaluate the expected, very faint dust emission from HVCs.

We have presented a new approach to address this issue by applying a GLM to evaluate the correlation of atomic neutral

hydrogen and FIR dust emission. The GLM offers the opportunity to attribute an individual dust emissivity to each spectral channel of the HI data. To regularise the fit, we introduced linear penalty terms for the GLM coefficients. The investigation of the residual FIR intensity shows that the GLM yields significantly lower residual emission than the standard approach. Furthermore, the distribution function of the residual is more symmetric, has fewer outliers, and the derived model is more robust to variations of the spatial area that is approximated.

After correcting for the foreground dust emission via this GLM, we analysed the potential dust content of the HVC. Using a line fit in the Bayesian framework, we derived that the dust emissivity at 857 GHz is $0.021 \text{ MJy sr}^{-1}/10^{20} \text{ cm}^{-2}$ at the 99.87% confidence level. This is more than an order of magnitude lower than typical ISM emissivities. This shows that even for this promising candidate with low kinetic temperatures, high brightness temperature, and head-tail structure, the detection of dust is not feasible via the correlation of dust and neutral gas.

6.2. Outlook

For the future search of dust in HVCs, we plan to apply a similar method to a larger sample of clouds and improve the signal-to-noise ratio by stacking them. Furthermore, spectroscopic observations of molecular gas tracers such as CO and OH (Allen et al. 2015) can help to shed light on the dust content of HVCs.

In another upcoming study, we will combine the different FIR frequencies as well as data on the gaseous content of the Milky Way such as HI and CO. Combined with a proper treatment of the CIB characteristics such the spatial covariance that needs to be considered for the analysis, we are confident that we will provide full-sky information on the dust emissivities, the X_{CO} conversion factor between molecular hydrogen and CO intensity ($\equiv N_{\text{H}_2}/W_{\text{CO}}$), and the CIB.

Acknowledgements. We thank the referee for providing very helpful feedback, especially for the role of the CIB anisotropies. We thank the Deutsche Forschungsgemeinschaft (DFG) for the project funding under grants KE757/7-1 to -3. Based on observations with the 100 m telescope of the MPIfR (Max-Planck-Institut für Radioastronomie) at Effelsberg. D.L. is a member of the Bonn-Cologne Graduate School of Physics and Astronomy (BCGS). L.F. is a member of the International Max Planck Research School (IMPRS) for Astronomy and Astrophysics at the Universities of Bonn and Cologne.

References

- Allen, R. J., Hogg, D. E., & Engelke, P. D. 2015, *AJ*, 149, 123
 Braun, R., & Burton, W. B. 2000, *A&A*, 354, 853
 Brüns, C., Kerp, J., & Pagels, A. 2001, *A&A*, 370, L26
 D’Agostini, G. 2003, *Rep. Progr. Phys.*, 66, 1383
 de Souza, R. S., Cameron, E., Killedar, M., et al. 2014, *Astronomy & Computing*, 12, 21
 Desert, F. X., Bazell, D., & Boulanger, F. 1988, *ApJ*, 334, 815
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
 Fox, A. J., Savage, B. D., Wakker, B. P., et al. 2004, *ApJ*, 602, 738
 Gillmon, K., Shull, J. M., Tumlinson, J., & Danforth, C. 2006, *ApJ*, 636, 891
 Goodman, J., & Weare, J. 2010, *Commun. Appl. Math. Comput. Sci.*, 5, 65
 Guillard, P., Boulanger, F., Pineau Des Forêts, G., & Appleton, P. N. 2009, *A&A*, 502, 515
 Jeffreys, H. 1946, *Proc. Roy. Soc. Lond. Ser. A*, 186, 453
 Kerp, J., Winkel, B., Ben Bekhti, N., Flöer, L., & Kalberla, P. M. W. 2011, *Astron. Nachr.*, 332, 637
 Madsen, H., & Thyregod, P. 2010, *Introduction to General and Generalized Linear Models*, Chapman & Hall/CRC Texts in Statistical Science (Taylor & Francis)
 Miville-Deschênes, M.-A., Boulanger, F., Reach, W. T., & Noriega-Crespo, A. 2005b, *ApJ*, 631, L57
 Muller, C. A., Oort, J. H., & Raimond, E. 1963, *Comptes Rendus Académie des Sciences, Paris*, 257, 1661

- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Machine Learning Res.*, 12, 2825
 Peek, J. E. G., Heiles, C., Putman, M. E., & Douglas, K. 2009, *ApJ*, 692, 827
 Picard, R. R., & Cook, R. D. 1984, *J. Am. Stat. Assoc.*, 79, 575
 Planck Collaboration XVIII. 2011, *A&A*, 536, A18
 Planck Collaboration XXIV. 2011, *A&A*, 536, A24
 Planck Collaboration Int. XVII. 2014, *A&A*, 566, A55
 Planck Collaboration XXX. 2014, *A&A*, 571, A30
 Planck Collaboration VIII. 2016, *A&A*, in press, DOI: 10.1051/0004-6361/201525820
 Putman, M. E., Peek, J. E. G., & Jounge, M. R. 2012, *ARA&A*, 50, 491
 Röhser, T., Kerp, J., Winkel, B., Boulanger, F., & Lagache, G. 2014, *A&A*, 564, A71
 Saul, D. R., Peek, J. E. G., & Putman, M. E. 2014, *MNRAS*, 441, 2266
 Tibshirani, R. 1996, *J. Roy. Stat. Soc. B (Methodological)*, 58, 267
 Wakker, B. P. 2001, *ApJS*, 136, 463
 Wakker, B. P., & Boulanger, F. 1986, *A&A*, 170, 84
 Wakker, B. P., & van Woerden, H. 1997, *ARA&A*, 35, 217
 Winkel, B., Kalberla, P. M. W., Kerp, J., & Flöer, L. 2010, *ApJS*, 188, 488
 Winkel, B., Kerp, J., Flöer, L., et al. 2016, *A&A*, 585, A41

Appendix A: Estimators in GLM reconstruction

We present the evaluation of the simulations for the spiky and flat GLM coefficients (Figs. A.1 and A.2, respectively). See Sect. 4.2 for a detailed description of the results and their implications.

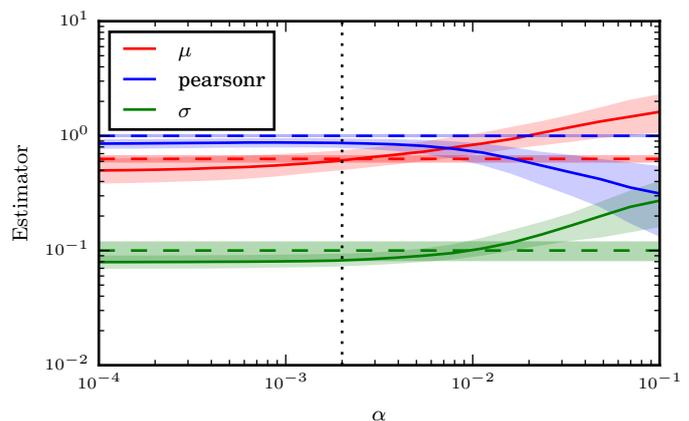


Fig. A.1. Evaluation of the reconstructed CIB mean μ , CIB fluctuation amplitude σ and Pearson’s r of input CIB image and reconstruction. This is based on the spiky input GLM coefficients (second panel from the top in Fig. 7). Dashed lines indicate the input, solid lines the reconstructed quantities. Contours correspond to 1σ uncertainties.

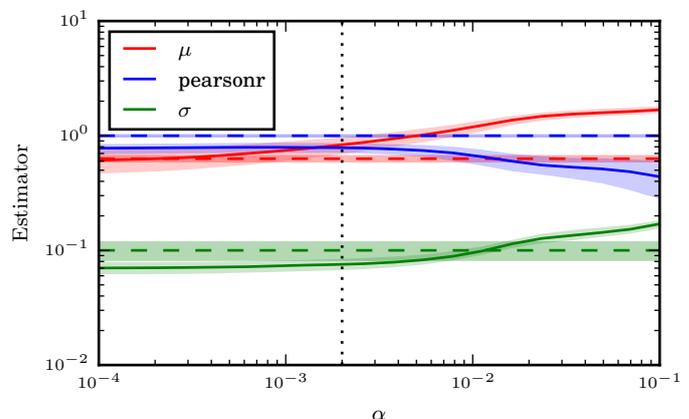


Fig. A.2. Same as Fig. A.1, but for a flat spectrum of GLM coefficients (bottom panel in Fig. 7).