

The ILIUM forward modelling algorithm for multivariate parameter estimation and its application to derive stellar parameters from Gaia spectrophotometry

C.A.L. Bailer-Jones

*Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany**

Accepted 2009 November 26. Received 2009 November 26; in original form 2009 October 19

ABSTRACT

I introduce an algorithm for estimating parameters from multidimensional data based on forward modelling. It performs an iterative local search to effectively achieve a nonlinear interpolation of a template grid. In contrast to many machine learning approaches it avoids fitting an inverse model and the problems associated with this. The algorithm makes explicit use of the sensitivities of the data to the parameters, with the goal of better treating parameters which only have a weak impact on the data. The forward modelling approach provides uncertainty (full covariance) estimates in the predicted parameters as well as a goodness-of-fit for observations, thus providing a simple means of identifying outliers. I demonstrate the algorithm, ILIUM, with the estimation of stellar astrophysical parameters (APs) from simulations of the low resolution spectrophotometry to be obtained by Gaia. The AP accuracy is competitive with that obtained by a support vector machine. For zero extinction stars covering a wide range of metallicity, surface gravity and temperature, ILIUM can estimate T_{eff} to an accuracy of 0.3% at $G=15$ and to 4% for (lower signal-to-noise ratio) spectra at $G=20$, the Gaia limiting magnitude (mean absolute errors are quoted). $[\text{Fe}/\text{H}]$ and $\log g$ can be estimated to accuracies of 0.1–0.4 dex for stars with $G \leq 18.5$, depending on the magnitude and what priors we can place on the APs. If extinction varies a priori over a wide range (0–10 mag) – which will be the case with Gaia because it is an all sky survey – then $\log g$ and $[\text{Fe}/\text{H}]$ can still be estimated to 0.3 and 0.5 dex respectively at $G=15$, but much poorer at $G=18.5$. T_{eff} and A_V can be estimated quite accurately (3–4% and 0.1–0.2 mag respectively at $G=15$), but there is a strong and ubiquitous degeneracy in these parameters which limits our ability to estimate either accurately at faint magnitudes. Using the forward model we can map these degeneracies (in advance), and thus provide a complete probability distribution over solutions. Additional information from the Gaia parallaxes, other surveys or suitable priors should help reduce these degeneracies.

Key words: surveys – methods: data analysis, statistical – techniques: spectroscopic – stars: fundamental parameters – ISM: extinction

1 INTRODUCTION

Inferring parameters from multidimensional data is a common task in astronomy, whether this be inference of cosmological parameters from CMB experiments, photometric-redshifts of galaxies or physical properties from stellar spectra. Important questions about the structure and evolution of stars and stellar populations require knowledge of abundances and ages, which must be obtained spectroscopically via the stellar atmospheric parameters effective temperature (T_{eff}), surface gravity ($\log g$) and iron-peak metallicity ($[\text{Fe}/\text{H}]$).

Numerous publications have presented methods for estimating such astrophysical parameters (APs) from spectra. The methods can be divided into two broad categories. The first, based on high resolution and high signal-to-noise ratio (SNR) spectra, makes use of specific line indices selected to be sensitive predominantly to the phenomena of interest. Examples include the detection of BHB stars via Calcium and Balmer lines (e.g. Brown 2003) and spectral type classification of M, L and T dwarfs via molecular band indices (e.g. Hawley et al. 2002). In each case these methods use a specific, identified phenomenon and so are generally limited to a narrow part of the AP space where the values of the other APs are reasonably well known. While simple and useful, these methods do not use all available information nor can they normally be used with low res-

* Email: calj@mpia.de

olution data, because the effects of individual APs cannot then be separated.

The second category of methods is global pattern recognition, which try to use the full set of available data. These are generally used when we want to estimate one or more APs over a wide range of parameter space. In such cases there is rarely a simple relation between a single feature and the physical quantity of interest. (An example is determination of surface gravity, where the relevant lines to use change with temperature.) We must therefore infer which features are relevant to which APs in which parts of the AP space. As this is generally a nonlinear, multidimensional problem, the standard approach is to use a machine learning algorithm to learn the mapping from the data space to the AP space based on labelled template spectra (spectra with known APs). Various models have been used in astronomy for a variety of problems, including (to name just a few): neural networks for stellar parameter estimation (e.g. Re Fiorentin et al. 2007) or photometric redshift estimation (e.g. Firth et al. 2003); support vector machines for quasar classification (e.g. Gao et al. 2003, Bailer-Jones et al. 2008) or galaxy morphology classification (e.g. Huertas-Company et al. 2008); classification trees for identifying cosmic rays in images (e.g. Salzberg et al. 1995); linear basis function projection methods (e.g. the method of Recio-Blanco et al. 2006 for spectral parameter estimation, which constructs the basis functions from model spectra). More examples of machine learning methods and their use in astronomy can be found in the volume edited by Bailer-Jones (2008).

Note that the first category of methods, line indices, is really just a special case of the second in which drastic feature selection has taken place to enable use of low (one or two) dimensional models. In both cases we must learn some relationship $AP = g'(Data)$ from a model or based on some labelled templates. But this is an *inverse* relation: more than one set of APs may fit a given set of data (e.g. a low extinction cool star or high extinction hot star could produce the same colours). Despite this non-uniqueness we nonetheless try and fit a unique model. This causes fitting problems which become more severe the lower the quality of the data (the lower the number of independent measures) and the larger the number of APs we want to estimate from it, and could lead to poor AP estimates or biases. In contrast, the forward mapping (or *generative model*), $Data = g(AP)$ is unique, because this is a causal, physical model (e.g. a stellar atmosphere and radiative transfer model of a spectrum). A further issue is that the model must learn the sensitivity of each input to each AP (and how noise affects this). Yet this information we in principle have already from the gradients of the generative model.

A pattern recognition method which tries to overcome some of these problems is *k*-nearest neighbours, which has also been applied to many problems in astronomy (e.g. Katz et al. 1998, Ball et al. 2008; plus extensions thereof such as kernel density estimation, e.g. Richards et al. 2004, or the method of Shkedy et al. 2007, which converts the distances into likelihoods and uses priors to create a full probabilistic solution.) In many ways this is the most natural way to solve the problem: we create a grid of labelled templates and find which are closest to our observation (perhaps smoothing over several neighbours). On the assumption that the APs vary smoothly with the data between the grid points, this may provide a good estimate. But for this to be accurate (and not too biased), the grid must be sufficiently dense that multiple grid points lie within the error ellipse of the observation (the covariance of these neighbours then provides a measure of the uncertainty in the estimated APs). If the SNR is high, the grid must therefore be very dense.

Moreover, as the number of APs increases, the required grid density grows exponentially with it: For a stellar parametrization problem with 5 APs we might need an average of 100 samples per AP, resulting in $100^5 = 10^{10}$ templates. There is also the issue of what distance metric to use. The covariance-weighted Mahalanobis distance is often used, but it ignores the sensitivities of the inputs. (If some inputs are sometimes dominated by irrelevant cosmic scatter, this will add unmodelled noise to the distance estimate.) This is a problem when we have a mix of APs, some of which have a large and others a small impact on the variance (“strong” and “weak” APs). If we just use the Mahalanobis distance we lose sensitivity to the weak APs.

We could overcome these problems if we did on-the-fly interpolation of the template grid to generate new templates as we need them. Running stellar models is far too time consuming for this, but also unnecessary because the generative model is smooth: We can instead fit a forward model to a low density grid of templates as an approximation to the generative model. As we shall see, possession of a forward model opens up opportunities not available to the inverse methods, such as direct uncertainty estimates and goodness-of-fit assessment of the solution.

In this article I introduce an algorithm for AP estimation based on this forward modelling idea and iterative interpolation. I will demonstrate it using simulations of low resolution spectra to be obtained from the Gaia mission (e.g. Lindegren et al. 2008).¹ Gaia will observe more than 10^9 stars down to 20^{th} magnitude over the whole Galaxy, stars which a priori span a very wide range in several APs. This includes the line-of-sight extinction parameter, A_V , which must be estimated accurately if we want to derive intrinsic stellar luminosities from the Gaia parallaxes. AP estimation (Bailer-Jones 2005) is therefore an integral part of the overall Gaia data processing and comprises one of the Coordination Units in the Gaia Data Processing and Analysis Consortium (DPAC) (Mignard & Drimmel 2007, O’Mullane et al. 2007).

I will now describe the basic algorithm (section 2). In section 3 I then introduce the simulated Gaia spectroscopy to which ILIUM is applied, with the results and discussion thereof presented in sections 4, 5 and 6. The latter section also reports on a strong and ubiquitous degeneracy between T_{eff} and A_V . I summarize and conclude in section 8. Additional plots, results and discussions can be found in a series of four Gaia technical notes (Bailer-Jones 2009a,b,c,d) available from <http://www.mpia.de/Gaia>.

2 THE ILIUM ALGORITHM

I outline the algorithm using the terminology of spectra and stellar astrophysical parameters, although it is quite general and applies to any multivariate data. Table 1 summarizes the notation. “Band” refers to a flux measurement in the spectrum. In general it could be a photometric band, a single pixel in the spectrum, or a function of many pixels.

2.1 Forward modelling

I will call the true relationship between the APs and the flux in a band *i* the *generative model*, $g_i(\phi)$. This provides the observed

¹ <http://www.rssd.esa.int/Gaia>

Table 1. Notation

I	number of bands (pixels in spectrum)
i	counter over band, $i = 1 \dots I$
p_i	photon counts in band i (\mathbf{p} is a spectrum)
J	number of APs (astrophysical parameters)
j	counter over AP, $j = 1 \dots J$
ϕ_j	AP j (ϕ is a set of APs)
s_{ij}	sensitivity of band i to AP j , $\frac{\partial p_i}{\partial \phi_j}$
\mathbf{S}	sensitivity matrix, $I \times J$ matrix with elements s_{ij}
$f_i(\phi)$	forward model for band i
n	iteration; e.g. $\phi(n)$ is the AP at iteration n

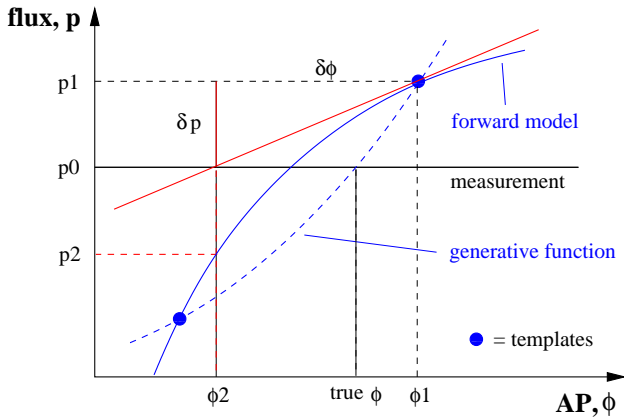


Figure 1. Sketch of the search method for one band ($I=1$) and one AP ($J=1$). The dashed blue curve is the (unknown) generative model, and the solid blue curve is the forward model (our approximation to the generative model) formed by fitting a function to the templates. (The difference between the two is exaggerated.) The straight red line shows the local linear approximation of the forward model (tangent at ϕ_1) used to calculate the first AP step. (In the case show the forward model could be inverted. But this is not generally the case, not even in one dimension if it has a turning point.)

spectrum for a given set of APs and thus encapsulates the underlying stellar model, radiative transfer, interstellar extinction, instrument model etc. (There is a separate function for each band, but for simplicity I will refer to it in the singular.) Generally we don't have an explicit function for this model, so it remains unknown. All we have for doing AP estimation is a discrete *template grid* of example spectra with known APs generated by the generative model. We approximate the generative model using a *forward model*, $f_i(\phi)$, which is a (nonlinear) parametrized fit to this grid and provides flux estimates at arbitrary APs, i.e. off the grid. (Forward models can be fit independently for each band.) Demanding the forward models to be continuous functions ensures we can also use them to calculate the sensitivities, which by definition are the gradients of the flux with respect to each AP. The forward model fitting is done just once for a given grid and is kept fixed when predicting APs. In other words it is a training procedure.

2.2 Core algorithm

The basic idea of ILIUM is to use the Newton-Raphson method to find that forward model-predicted spectrum (and associated APs) which best fits the observed spectrum.

In detail, the algorithm is as follows (Fig. 1). Consider first a single AP and single band. The measurement is $p(0)$ and we want

to estimate its AP. The forward model, $\hat{p} = f(\phi)$, has been fit and remains fixed. The procedure is as follows (n is the iteration number)

(i) Initialize: find nearest grid neighbour to $p(0)$, i.e. the one which minimizes the sum-of-squares residual $\delta \mathbf{p}^T \delta \mathbf{p}$. Call this $[p(1), \phi(1)]$. $\phi(1)$ is the initial AP estimate.

(ii) Use the forward model to calculate the local sensitivities, $\frac{\partial p}{\partial \phi}$, at the current AP estimate.

(iii) Calculate the discrepancy (residual) between the predicted flux and the measured flux, $\delta p(n) = p(n) - p(0)$.

(iv) Estimate the AP offset as $\delta \phi(n) = \left(\frac{\partial \phi}{\partial p} \right)_{\phi(n)} \times \delta p(n)$, i.e. a Taylor expansion truncated to the linear term. (Note that this partial derivative is the reciprocal sensitivity.)

(v) Make a step in AP space, $\phi(n+1) = \phi(n) - \delta \phi(n)$, toward the better estimate. This is the new AP prediction.

(vi) Use the forward model to predict the corresponding (off-grid) flux, $p(n+1)$

(vii) Iterate steps ii–vi until convergence is achieved or a stop is imposed.

The algorithm is basically minimizing $|\delta p|$. At each iteration we obtain an estimate of the APs (step v) and the corresponding spectrum (step vi). Convergence could be defined in several ways, e.g. when changes in the spectrum or the APs (or their rate of change) drop below some threshold. Alternatively we could simply stop after some fixed number of iterations. There is no guarantee of convergence. For example, if the AP steps were sufficiently large to move to a part of the function with a sensitivity of the opposite sign, then the model could diverge or get stuck in a limit cycle. Likewise, if initialized too far from the true solution the algorithm could become stuck in a local minimum far from the true solution. For this, and other reasons, the algorithm in practice has some additional features (discussed in section 2.5). Also,

2.3 Generalization to multiple APs and bands

In general we have several bands and several APs. The flux perturbation due to small changes in the APs is then

$$\delta \mathbf{p} = \mathbf{S} \delta \phi \quad (1)$$

where \mathbf{S} is the $I \times J$ sensitivity matrix with elements $s_{ij} = \partial p_i / \partial \phi_j$. Note that $I > J$. Multiplying this equation on the left by $(\mathbf{S}^T \mathbf{S})^{-1} \mathbf{S}^T$ gives

$$\delta \phi = (\mathbf{S}^T \mathbf{S})^{-1} \mathbf{S}^T \delta \mathbf{p} \quad (2)$$

so the AP update equation (step v in the algorithm) becomes

$$\phi(n+1) = \phi(n) - (\mathbf{S}^T \mathbf{S})^{-1} \mathbf{S}^T \delta \mathbf{p}(n) \quad (3)$$

The I forward models are now functions of J variables, and this turns out to be a critical matter.

2.4 The forward model

The core algorithm just described can make use of any form for the forward model, on the condition that it provides values of the function and its first derivatives for arbitrary values of the APs.

The most obvious forward model would be a multidimensional, nonlinear regression of the form $\hat{p} = f(\phi)$, which in principle works for any number of APs. However, I found that it was difficult to get a model which simultaneously fits both T_{eff} , a strong AP, and $\log g$, a weak AP to sufficient accuracy. *Strong* here means




Figure 2. Schematic diagram of a two-component forward model. Note the stronger variation in the flux in the direction of the strong AP (the contrast is typically much larger for the problems considered in this article).




Figure 3. Schematic diagram of the grid of AP values, $k = 1 \dots 20$, $l = 1 \dots 7$. The solid (red) points denote those used to fit the weak component of the forward model for one value of the strong AP.

that it explains much of the variance in the flux data, i.e. is a strong predictor of the flux. *Weak* is a relative term, indicating that the AP explains much less of the variance. The reason for poor fits over the weak APs is that the model is fit by minimizing a single objective function, namely the error in reproducing the flux. As the weak AP has very little impact on the flux, its influence has little impact on the error, so the model optimization does little to produce a good fit over this AP.

To overcome this problem I use a two-component forward model to separately fit the strong and weak APs. Consider the case of a single strong AP and a single weak AP (this is generalized later). The strong component is a 1D nonlinear function of the strong AP which is fit by marginalizing over the weak AP. This fits most of the flux variation. Then, at each discrete value of the strong AP in the grid, we fit the residual flux as a function of the weak AP; these are the second components (also 1D). They provide a flux increment dependent on the weak AP, which is added to the flux predicted given the strong AP. A schematic illustration of such a two-component forward model is illustrated in Fig. 2.

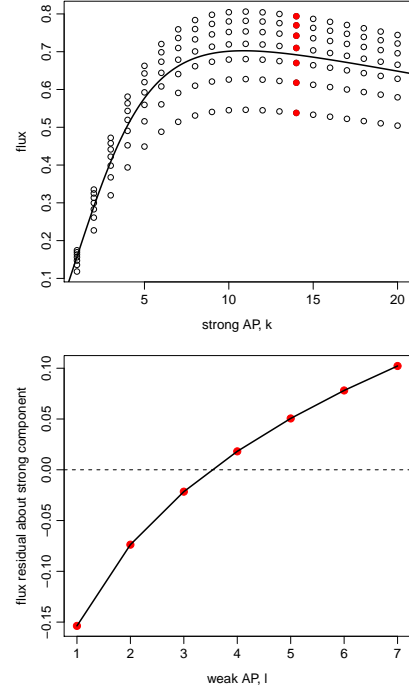


Figure 4. Schematic illustration of the two components of the forward model fit over all the strong and weak points in the grid in Fig. 3 Top: the fit over the strong APs. Bottom: one of the fits over the weak APs (the solid/red points in the top panel at $k=14$).

To be more precise the model is fit as follows. Let ϕ^S denote a strong AP, ϕ^W a weak AP and $f_i(\phi^S, \phi^W)$ the complete (2D) forward model for band i . Let subscripts k and l denote specific values in the grid of the strong and weak APs respectively (see Fig. 3). We model the flux at an arbitrary AP point as

$$f_i(\phi^S, \phi^W) = f_i^S(\phi^S) + f_{i,k}^W(\phi^W; \phi^S = \phi_k^S) \quad (4)$$

where both f_i^S and $f_{i,k}^W$ are 1D functions. $f_i^S(\phi^S)$ is the single strong component (for band i). It is a fit to the average value of ϕ^W at each ϕ^S (the curve in the top panel of Fig. 4). $f_{i,k}^W$ is the k^{th} weak component, which is a fit with respect to the weak AP with the strong AP fixed at ϕ_k^S (e.g. the solid/red points in Fig. 3), i.e. it is fit to the residuals

$$\{p(\phi^S = \phi_k^S) - \overline{p(\phi^S = \phi_k^S)}\}$$

as illustrated in in the bottom panel of Fig. 4.

This fitting approach clearly requires us to have a “semi-regular” grid: one which has a range of values of the weak AP for each value of the strong AP. (This requirement is easily fulfilled when using synthetic grids.) The number of weak components in the model is equal to the number of unique values of ϕ^S , 20 in this schematic case. If we have more than one strong or weak component then we raise the dimensionality of the strong or weak component (see section 6).

Applying the forward model is easy. Given (ϕ^S, ϕ^W) we

- (i) evaluate f_i^S , the strong component;
- (ii) find the nearest value, ϕ_k^S , in the grid to ϕ^S , i.e. identify the closest weak component;
- (iii) evaluate $f_{i,k}^W$, the increment from the weak component;

(iv) sum the two components, $f_i^S + f_{i,k}^W$, to give the forward model prediction.

Although the weak component we use changes discontinuously as ϕ^S varies, the weak component is only specifying an *increment* to the strong component fit. As both components are smooth in their respective APs the combined function is also smooth along any direction parallel to the AP axes. It is not smooth along arbitrary directions, but this is unimportant because explicit calculations with the forward model (e.g. of the sensitivities) are only carried out parallel to the AP axes.

2.5 Practical algorithm

The practical realities of working with real (noisy) data mean that the basic algorithm should be extended in order to make it more robust. These I now describe together with other implementation aspects used for the experiments described later. For the purposes of this paper the algorithm has been implemented in \mathbb{R}^2 . A Java implementation is in progress, which will be necessary for larger scale applications.

2.5.1 Standardized variables

The spectral variables are observed photon counts (to within an irrelevant constant factor). The APs are all on logarithmic scales: A_V in magnitudes, $\log g$ and $[\text{Fe}/\text{H}]$ in dex, and $\log(T_{\text{eff}})$. In order to bring each variable to the same level I standardize the flux in each band and each AP (linearly scale each to have zero mean and unit variance). If there were correlations in the spectra these could be removed by “spherizing” (“prewhitening”), a covariate generalization of standardization which gives the data unit diagonal covariance (e.g. Bishop 2006).

2.5.2 Forward model functions

The strong and weak components of the forward model are fit using smoothing splines (e.g. Hastie et al. 2001). Conventional cubic splines have the drawback that one must control their complexity (smoothness) using the number and position of the knots. Smoothing splines circumvent this problem by setting a knot at every point (which would overfit the data) and then applying a smoothing penalty which is controlled by specifying the effective degrees-of-freedom (dof). I set this by trial and error via inspection of the resulting fits. For the 2D problems TG (T_{eff} and $\log g$) and TM (T_{eff} and $[\text{Fe}/\text{H}]$) described in sections 4 and 5, both the strong and the weak model splines are 1D. The strong model uses $\text{dof} = n_{T_{\text{eff}}}/2 = 16.5$ where $n_{T_{\text{eff}}}$ is the number of unique T_{eff} points. As the maximum number of $\log g$ points is 10 (for the training data), and because the variation with $\log g$ is quite smooth, I set the dof for these fits to be 4. However, many of the T_{eff} values in the training grid have fewer $\log g$ or $[\text{Fe}/\text{H}]$ points: To avoid overfitting, if $n_{\log g}(T_{\text{eff}}) \leq 4$ then a linear fit is used. If $n_{\log g}(T_{\text{eff}}) = 1$, then no fit is performed and this weak component of the forward model is zero. (Practical aspects of higher dimensional forward models are described in section 6.)

The forward model is always fit to noise-free data. It is well known (and the author’s experience) that inverse modelling methods such as ANNs and SVMs perform best when trained on

data with a similar noise level as the target data (e.g. Snider et al. 2001). This precipitates the need for multiple models when used on survey data with a range of SNRs, something which is not an issue for ILIUM.

2.5.3 Sensitivity estimation

If the forward model is a simple analytic function then it may have analytical first derivatives which can be used to calculate the sensitivities. But in the general case we can use the method of first differences

$$\left. \frac{\partial p}{\partial \phi_j} \right|_{\phi} \simeq \frac{f(\phi + \delta \phi_j) - f(\phi - \delta \phi_j)}{2 \delta \phi_j} \quad (5)$$

I select $\delta \phi_j$ to be slightly smaller than the maximum precision a priori possible in an AP. As the forward model must be smooth at this resolution, the first difference approximation is sufficiently accurate. For the examples shown later, I choose $\delta \phi_j$ to be 0.05 dex for $\log g$ and $[\text{Fe}/\text{H}]$, 0.0005 for $\log(T_{\text{eff}})$ (0.1% for T_{eff}) and 0.03 mag for A_V .

2.5.4 Lower limit on sensitivities (singularity avoidance)

Given that the AP updates depend upon the inverse of the sensitivity matrix (equation 2), it is prudent to prevent the sensitivities being too small in order to avoid $\mathbf{S}^T \mathbf{S}$ being singular. For this reason, a lower limit is placed on the absolute value of each sensitivity, s_{ij} , of 0.001 (with p and ϕ in standardized units). For the examples shown later, this limit rarely had to be applied in practice and had negligible impact on the results. To avoid singularity $\mathbf{S}^T \mathbf{S}$ must also have a rank of at least J , i.e. to estimate J APs we need at least J independent measures in the spectrum.

2.5.5 AP update contribution clipping

Equation 2 can be written

$$\phi(n+1) = \phi(n) - \mathbf{M} \delta \mathbf{p}(n) \quad (6)$$

where

$$\mathbf{M} = (\mathbf{S}^T \mathbf{S})^{-1} \mathbf{S}^T \quad (7)$$

is a $J \times I$ matrix. Equation 6 gives J update equations, one for each AP. The update for AP j can be written as the dot product of two vectors, the j^{th} row of \mathbf{M} , \mathbf{m}_j with $\mathbf{p}(n)$, i.e.

$$\begin{aligned} \phi_j(n+1) &= \phi_j(n) - \mathbf{m}_j \delta \mathbf{p}(n) \\ &= \phi_j(n) - \sum_i m_{ij} \delta p_i(n) \\ &= \phi_j(n) - \sum_i u_{ij}(n) \end{aligned} \quad (8)$$

which defines u_{ij} . Thus we see that the update to AP j is a sum over I individual updates, which we can view as an update “spectrum”. If we inspect these updates (see Bailer-Jones 2009a) we see that on occasion some are much larger than the others. This dominance of the update by just one of a few elements is undesirable, because they may be affected by noise ($\delta \mathbf{p}(n)$ is a noisy measurement). For this reason, I clip outliers in this spectrum. (It is valid to compare the updates for different bands, because we work with standardized fluxes.) To be robust, I set an upper (lower) limit which is a multiple

² <http://www.r-project.org>

c of the median of those points above (below) the median. Using the notation $\theta()$ to denote median, the limits are

$$\begin{aligned} u_{\text{upper}} &= \theta(u_i) + c[\theta(u_i > \theta(u_i)) - \theta(u_i)] \\ u_{\text{lower}} &= \theta(u_i) + c[\theta(u_i < \theta(u_i)) - \theta(u_i)] \end{aligned} \quad (9)$$

I somewhat arbitrarily set $c = 10$ so as to be relatively conservative in clipping.

2.5.6 Upper limit on AP update step size

The AP steps at each iteration (equation 2) could be very large. This is undesirable, because the updates are based on a local *linear* approximation to the generative model. The code therefore imposes upper limits on the AP updates, corresponding to steps no larger than about 2.0 dex in $\log g$ and $[\text{Fe}/\text{H}]$, 0.04 in $\log(\text{T}_{\text{eff}})$ (10% in T_{eff}) and 0.3 mag in A_V . A larger step size is permitted for the weaker APs because the initial nearest neighbour offset can be quite incorrect. These limits are imposed more often for noisy data, but still relatively rarely.

2.5.7 Limit AP extrapolation

We do not expect the forward model to make good predictions beyond the AP extremes of the grid, so I set upper and lower limits on the AP estimates which ILIUM can provide. These are set as e times the range of each AP, i.e.

$$\begin{aligned} \text{upper limit} &= \max \phi_j + e(\max \phi_j - \min \phi_j) \\ \text{lower limit} &= \min \phi_j - e(\max \phi_j - \min \phi_j) \end{aligned} \quad (10)$$

I set $e = 0.1$.

2.5.8 Stopping criterion

The algorithm is simply run for a fixed number of iterations (20). We often observe good natural convergence, so a more sophisticated stopping criterion is not applied at this time, although it may be important on high variance data sets (low SNR or more APs).

2.6 Performance statistics

The model performance is assessed via the AP residuals (estimated minus true, $\delta\phi$) on an evaluation data set. I report three statistics: (1) the root-mean-square (RMS) error, which I abbreviate with ϵ_{rms} ; (2) the mean absolute error, $|\delta\phi|$, abbreviated as ϵ_{mae} ; (3) the mean residual, $\bar{\delta\phi}$, a measure of the systematic error, abbreviated as ϵ_{sys} . (Note that as (1) and (2) are statistics with respect to the true values they also include any systematic errors.) I mostly use ϵ_{mae} rather than RMS because the former is more robust. If the residuals had a Gaussian distribution then the RMS would equal the Gaussian 1σ which is $\sqrt{\pi/2} = 1.25$ times larger than ϵ_{mae} . But usually there are outliers which increase the RMS significantly beyond this.

2.7 Uncertainty estimates

If vectors \mathbf{y} and \mathbf{x} are related by a transformation $\mathbf{y} = \mathbf{A}\mathbf{x}$ then a standard result of matrix algebra is that the covariance of \mathbf{y} is $\mathbf{C}_{\mathbf{y}} = \mathbf{A}\mathbf{C}_{\mathbf{x}}\mathbf{A}^T$ where $\mathbf{C}_{\mathbf{x}}$ is the covariance of \mathbf{x} . Applying this to equation 2 gives us an expression for the covariance in the APs

$$\mathbf{C}_{\phi} = (\mathbf{S}^T\mathbf{S})^{-1}\mathbf{S}^T\mathbf{C}_p\mathbf{S}(\mathbf{S}^T\mathbf{S})^{-1} \quad (11)$$

as a function of the sensitivity (calculated at the estimated APs) and the covariance in the measured photometry, \mathbf{C}_p . (This equation assumes that ILIUM provides unbiased AP estimates and that the sensitivities have zero covariance. It can also be written $\mathbf{C}_{\phi} = \mathbf{M}\mathbf{C}_p\mathbf{M}^T$ where \mathbf{M} is the update matrix introduced in equation 7.) \mathbf{C}_p can be estimated from a photometric error model, and will be diagonal if the photometric errors in the bands are independent. Even in this case \mathbf{C}_{ϕ} is generally non-diagonal: the AP estimates are correlated on account of the sensitivities.

Because we have a forward model we can calculate a goodness-of-fit (GoF) for any estimate of the APs. Here I simply use the reduced- χ^2 to measure the difference between the observed spectrum and predicted spectrum³

$$\text{GoF} = \frac{1}{I-1} \sum_{i=1}^{i=I} \left(\frac{p_i - \hat{p}_i}{\sigma_{p_i}} \right)^2 \quad (12)$$

where $\hat{p}_i = f_i(\phi_j)$ is the forward model prediction and $\{\sigma_{p_i}^2\} = \text{diag}(\mathbf{C}_p)$ is the expected photometric noise. (Despite the name, a larger value refers to a poorer fit!) As the GoF can be measured without knowing the true APs, it can be used for detection of poor solutions or outliers.

Conventional methods of AP estimation via direct inverse modelling (e.g. with SVMs or ANNs) do not naturally provide uncertainty estimates and must usually resort to time-intensive sampling methods, such as resampling the measured spectrum according to its estimated covariance. They cannot provide a GoF at all because they lack a forward model.

2.8 Signal-to-noise weighted AP updates

The update equation (2) only takes into account the sensitivity of the bands, not their SNR. However, even if a band is very sensitive to an AP in principle, if its measurement is very noisy then it is less useful. We could accommodate this by including a factor proportional to \mathbf{C}_p^{-1} into equation 2 which would down-weight noisier measurements. Preliminary results using this on the TG problem (see section 3.2) show it actually degrades performance at $G=15$, but gives some improvement at $G=18.5$ (Bailer-Jones 2009c).

3 ASSESSING THE ALGORITHM

3.1 Gaia simulations

To illustrate ILIUM I apply it to estimate stellar APs from simulated Gaia stellar spectra and thereby also make preliminary predictions of the mission performance. Gaia will observe all of its targets with two low-dispersion slitless prism spectrographs, together covering the wavelength range from 350–1050 nm. (These are creatively called “BP” for blue photometer and “RP” for red photometer.) The dispersion varies from 3 nm/pixel at the blue end to 30 nm/pixel at the red end (Brown 2006). The blue and the red spectra are each sampled with 60 pixels, but as the line-spread-function of the spectrograph is much broader, these samples are not independent. After removing low SNR regions of the modelled spectra, I retain 34 pixels in BP covering 338–634 nm and 34 pixels in RP covering 667–1035 nm. This is a slightly narrower range (and 18 pixels fewer) than the one adopted by Bailer-Jones et al. (2008) for quasar classification with similar spectra.

³ I use $I-1$ degrees of freedom rather than I because all the spectra have a common G magnitude, so the bands are not all strictly independent.




Figure 5. The $T_{\text{eff}}\text{-log } g$ grid of the data used in the experiments.

Extensive libraries of BP/RP spectra have been simulated by the Gaia DPAC using the GOG (Gaia Object Generator; Luri et al. (2005), Isasi 2009) instrument model and libraries of input spectra.⁴ Here I use the Basel (Lejeune et al. 1997) and Marcs (e.g. Gustafsson et al. 2008) stellar libraries. The former (as used here) includes 17 T_{eff} values from 8000–15 000 K with non-uniform spacing and the latter 17 T_{eff} values from 4000–8000 K in uniform steps of 250 K. (Together there are 33 unique T_{eff} values because 8000 K is in both.) Together the grids span $\log g$ values from -0.5 to 5.0 dex in steps of 0.5 dex, although the grid is incomplete for astrophysical reasons (Fig. 5). $[\text{Fe}/\text{H}]$ ranges from -4.0 dex to $+1.0$ dex with 13 discrete values for the cooler stars (with $T_{\text{eff}} \leq 8000$ K): The spacing is 0.25 dex from $+1$ to -1 , followed by points at -1.5 , -2.0 , -3.0 and -4.0 . Not all $[\text{Fe}/\text{H}]$ are present at all $T_{\text{eff}}\text{-log } g$ combinations. Each star has been simulated at one of ten values of the interstellar extinction, $A_V \in \{0, 0.1, 0.5, 1, 2, 3, 4, 5, 8, 10\}$ with $R_V = 3.1$. (Note that A_V is the extinction parameter defined by Cardelli et al. 1989. It is not the extinction in the V band.) The Marcs library additionally shows variation in the alpha element abundances, $[\alpha/\text{Fe}]$, representing five values from $0.0\text{--}0.4$ dex in 0.1 dex steps. Hence this combined library shows variance in five APs, T_{eff} , A_V , $\log g$, $[\text{Fe}/\text{H}]$, $[\alpha/\text{Fe}]$, the first two of which are “strong” and the latter three “weak”. The total number of spectra is 46 310. ILIUM will be used to estimate the first four APs; $[\alpha/\text{Fe}]$ will be ignored and contributes cosmic scatter.

GOG simulates the number of photoelectrons (“counts”) in the spectral bands. (While these will be calibrated in physical flux units before being published to the community, the classification work by DPAC is currently done in photoelectron space.) The variance in the spectra due to the four APs of interest is demonstrated in the example spectra plotted in Figs. 6, 7, 8 and 9. The first plot shows the variance due to T_{eff} only. The break between the BP and RP instruments around 660 nm is clear, as is the highly variable dispersion. The lower counts in RP immediately to the right of the break compared to BP is primarily due to the higher dispersion (fewer photons per band). In the other plots two APs are varied while the other two are held constant. Fig. 7 demonstrates the strong impact of both T_{eff} and A_V variations. (As I will discuss in section 6.4,

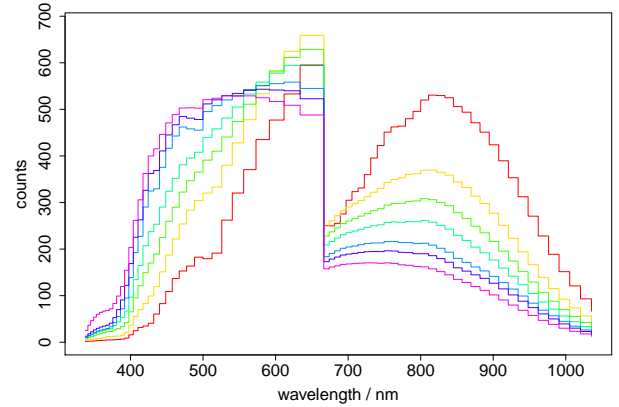


Figure 6. Noise-free Gaia spectra for solar metallicity dwarfs at zero extinction with $T_{\text{eff}} = \{4000, 5000, 6000, 7000, 8500, 10000, 15000\}$ K, increasing monotonically from bottom (red) to top (violet) at long wavelengths. They are composed of two spectra (BP and RP) to the left and right of about 660 nm. The ordinate is in units of photoelectrons (to within some constant multiple), not energy flux.




Figure 7. Noise-free Gaia spectra at a range of T_{eff} (as in Fig. 6) and A_V ($0.0, 0.1, 0.5, 1, 2, 3, 4, 5, 8, 10$) ranging from 0.0 mag (lowest line at long wavelengths; in red) to 10.0 mag (highest line at long wavelengths; in violet). Each temperature block has been offset by 1200 counts for clarity (the zero levels are shown by the dashed lines). $\log g = 4.0$ dex, $[\text{Fe}/\text{H}] = 0.0$ dex and $[\alpha/\text{Fe}] = 0.0$ dex in all cases.

⁴ For Gaia pundits: I use the CU8 cycle 3 simulations of the nominal (discrete) libraries (Sordo & Vallenari 2008)




Figure 8. Noise-free Gaia spectra at a range of T_{eff} and $\log g$ ($-0.5, 0.5, 2, 3, 4, 5$), with the lowest gravity (red) forming the lowest curve at the red end of the spectrum and the highest gravity (violet) the highest. Note that not all gravities are present at all T_{eff} due to the limitations of reality. Each temperature block has been offset by 600 counts for clarity (the zero levels are shown by the dashed lines) $A_V = 0.0$ mag, $[\text{Fe}/\text{H}] = 0.0$ dex and $[\alpha/\text{Fe}] = 0.0$ dex in all cases.

these two APs are highly degenerate in these data.) The third and fourth figures clearly demonstrate why $\log g$ and $[\text{Fe}/\text{H}]$ are weak parameters: they have little impact on the spectra compared to the variance due to T_{eff} or A_V . In particular, at high temperatures the spectra show essentially no sensitivity to $[\text{Fe}/\text{H}]$.

The AP estimation accuracy is, of course, a strong function of the SNR in the spectrum. As Gaia has a fixed sky scanning law, the SNR depends on the source magnitude and the number of observations (because the individual observations are combined into a single end-of-mission spectrum). I adopt here 72 observations for all spectra (the mean number of observations per source for a 5-year mission), and report results as a function of just the magnitude. The simulator noise model takes into account the source, background and background-subtraction photon (Poisson) noise as well as the CCD readout noise. At present errors due to the combination of spectra, charge-transfer inefficiency and CCD radiation damage are not explicitly included. There is instead a factor to account for general processing and calibration errors. All of these noise terms are combined into a zero mean Gaussian model for each pixel, the standard deviation of which is a function of the G-band apparent magnitude. (The G-band is the filterless band – defined by the mirror

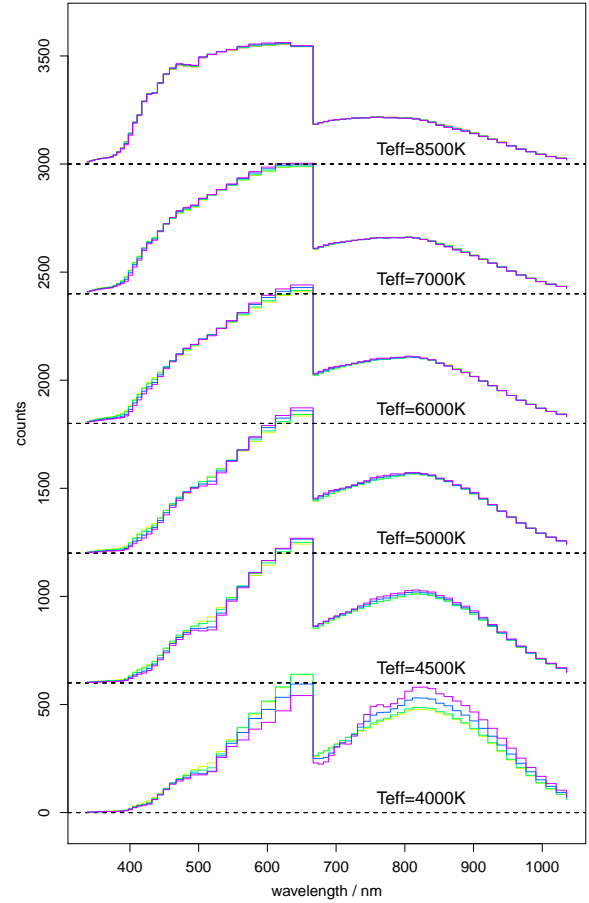


Figure 9. Noise-free Gaia spectra at a range of T_{eff} and $[\text{Fe}/\text{H}]$ ($-3, -2, -1, 0, +0.5$), with the lowest metallicity (red) forming the lowest curve at the red end of the spectrum and the highest metallicity (violet) the highest. Note that not all metallicities are present at all T_{eff} due to limitations of the simulated libraries. Each temperature block has been offset by 600 counts for clarity (the zero levels are shown by the dashed lines) $A_V = 0.0$ mag, $\log g = 4.0$ dex and $[\alpha/\text{Fe}] = 0.0$ dex in all cases.

and CCD response – in which the Gaia astrometry is obtained, covering a range similar to BP/RP.) This defines a “sigma spectrum” for each star from which I generate random numbers in order to simulate noisy spectra at $G=15, 18.5$ and 20 . The resulting SNR is a strong function of wavelength and the specific source, and is summarized in Fig. 10.

3.2 Data sets and forward model fitting

In the following sections I will apply Gaia to four distinct problems according to the APs we are trying to determine. In each case the forward model is fit to a grid varying in at least those APs, and in some cases the grid also shows variance in another AP (which therefore acts to provide additional “cosmic scatter”). For each case the ILIUM forward model is fit (trained) using the full, noise-free data set over the complete range of the represented APs. The corresponding noisy data set (at $G=15, 18.5$ or 20) is split randomly into two equal halves: one half is used for initialization (selecting the nearest neighbour) and ILIUM is applied to the other half (or a random subset of it where indicated below) on which the perfor-




Figure 10. The median signal-to-noise ratio (SNR) (solid line) and 0.1 and 0.9 quartiles (dashed lines) across the set of zero extinction dwarf stars for $G=18.5$ (red/thicker lines) and $G=20.0$ (blue/thinner lines). Compared to the $G=18.5$ curve, the SNR at $G=15$ is 8–6 times larger between 400 and 660 nm and 7–12 times larger between 660 and 1000 nm.

mance is evaluated.⁵ (The legitimacy of this procedure for evaluating performance is discussed in appendix A.) In addition to reporting global results (over the full AP ranges present in the training set) I also measure performance on subsets of the evaluation data set, in particular the $[\text{Fe}/\text{H}]$ just for cool stars (the ILIUM forward model is *not* refit). The problems and their grids are as follows (the name indicates the APs being determined: Temperature, Gravity, Metallicity and/or Extinction)

- TG: Estimation of $T_{\text{eff}} + \log g$ (1 strong and 1 weak AP), for stars with $A_V = 0$ and $[\text{Fe}/\text{H}] = 0$, some 274 stars. I also build a second TG model (TG-allmet) which is trained and evaluated with the full $[\text{Fe}/\text{H}]$ range in the grid (-4 to $+1$ dex). This contains 4361 stars, of which a quarter are used in the evaluation set.

- TM: Estimation of $T_{\text{eff}} + [\text{Fe}/\text{H}]$ (1 strong and 1 weak AP), for stars with $A_V = 0$ and either $\log g \in \{4.0, 4.5, 5.0\}$ (TM-dwarfs; 1716 stars), or $\log g \in \{1.0, 1.5, 2.0, 2.5, 3.0\}$, (TM-giants; 1882 stars). I also build a third model (TM-allgrav) which is trained and evaluated with the full $\log g$ range in the grid (-0.5 to $+5$ dex). This contains 4361 stars (it’s the same grid as used for TG-allmet of course), of which a quarter are used in the evaluation set.

- TAG: Estimation of $(T_{\text{eff}}, A_V) + \log g$ (2 strong and 1 weak AP), for stars with $[\text{Fe}/\text{H}] = 0$. This has 2740 stars, of which a random selection of 1000 is used for evaluation.

- TAM: Estimation of $(T_{\text{eff}}, A_V) + [\text{Fe}/\text{H}]$ (2 strong and 1 weak AP), for dwarfs with $\log g \in \{4.0, 4.5, 5.0\}$. This has 17 160 stars of which a random selection of 1000 is used for evaluation.

- TGM: Estimation of $T_{\text{eff}} + (\log g, [\text{Fe}/\text{H}])$ (1 strong and 2

⁵ These target spectra must first be normalized to have the same counts level as those on which ILIUM was trained. For this I just use the G magnitude to scale the counts. However, as the G -band is not identical to the BP/RP band adopted, this gives rise to a normalization offset between spectra even for a common G magnitude. For example, over the TG grid the integrated BP/RP counts varies by up to 10%, dependent primarily on T_{eff} . A better normalization might be “area normalization”, i.e. dividing the counts in each band by the sum over all bands for that spectrum. My G -band normalization is in principle conservative, as it mimics a small calibration bias.

weak APs, for stars with $A_V = 0$. This has 4361 stars of which a random selection of 1000 is used for evaluation.

In appendix B the ILIUM results are compared with results from an SVM on some of these problems.

4 APPLICATION TO THE $T_{\text{eff}} + \log g$ PROBLEM (TG)

First the forward model is fit for each band to the $T_{\text{eff}} - \log g$ grid shown in Fig. 5. As a reminder, each forward model comprises the 1D function over T_{eff} (the strong component) and 33 1D functions in $\log g$ (the weak components), as described in section 2.4. All of these are smoothing splines (section 2.5.2). Figs. 11 and 12 show the forward model fit for 12 bands at cuts of constant $\log g$ and T_{eff} respectively. The fits are good: they show the degree of smoothness we would expect for these data plus a robust extrapolation. If we compare the flux scales between Figs. 11 and 12 (these are standardized variables) we see how small the flux variation is as $\log g$ varies over its full range compared to T_{eff} : this is what it means to be a weak AP. Note also the small discontinuity in some of the bands in Fig. 11 at 8000 K ($\log(T_{\text{eff}}) = 3.903$) where the Marcs and Basel libraries join.

Having trained ILIUM I apply it to the evaluation data set at $G=15$. Fig. 13 shows five examples of how the AP estimates evolve. The first iteration is the nearest neighbour initialization; the final is the adopted AP estimate. The red (horizontal) lines show the true APs. Looking at many examples we see a range of convergence behaviours. Sometimes convergence is rapid, for other stars it takes longer. Sometimes it is smooth, other times not. It can be quite different for the two APs for a given star and depends also on the specific spectrum (which is noisy). Sometimes the nearest neighbour estimate is the correct one, and ILIUM may actually iterate away from this and converge on a different value. Convergence (on something) is almost always achieved on this problem, even though there is nothing adaptive in the algorithm. This is an encouraging property. Limit cycles are also seen in a handful of cases, but with negligible amplitudes on this problem (high SNR). For this specific problem, the extrapolation limits on the AP estimations (section 2.5.7) and the limits on the AP updates at each step (section 2.5.6) never had to be enforced by the algorithm.

Fig. 14 plots the ILIUM residuals (estimated minus true APs) on the 137 stars in the evaluation set; the statistics are summarized in line 1 of Table 2. We see that the APs can be estimated very accurately (no significant systematic error, ϵ_{sys}) and very precisely (low scatter, ϵ_{mae} or ϵ_{rms}). Applying the same model to noisier spectra obviously increases the errors (lines 2 and 3 in the table), but at $G=18.5$ the mean absolute errors are still an acceptable 1% in T_{eff} and 0.35 dex in $\log g$. At $G=20$ $\log g$ cannot be estimated accurately enough to reliably distinguish dwarfs and giants, although T_{eff} is still okay with an expected error of 260 K at 6000 K, for example.

This evaluation set is relatively small so the error statistics are subject to variation. Applying ILIUM to an ensemble of randomly selected evaluation sets (at $G=18.5$), we see that the error statistics vary by 5–10% (inter-quartile range) of their mean. The reported values are therefore reasonably representative.

The performance of an SVM on these TG problems is given in Table B1. For both APs and all three magnitudes ILIUM is similar to or significantly better than SVM. (As there is some variance in the results from both methods I only consider the performance significantly different if the better one has a ϵ_{mae} at least 25% smaller.) So in this limited variance problem, at least, the forward modelling approach improves performance.

Table 2. Summary statistics (defined in section 2.6) of the ILIUM performance on various experiments. The first column indicates the model and grid used (defined in section 3.2). The second column defines the magnitude and composition of the evaluation data set: F means the full AP ranges (as present in the training set); L means only stars with $T_{\text{eff}} \leq 7000$ K. “<” in the ϵ_{sys} column indicates that the systematic error has a magnitude less than twice the standard error in the mean ($|\epsilon_{\text{sys}}| \leq 2 \times 1.25 \epsilon_{\text{mae}} / \sqrt{N}$), i.e. is statistically insignificant. (Quite a few not so marked are only marginally significant.) The units of the error statistics are logarithmic for all variables (dex for all except A_V which is in magnitudes) and are on the proper variables (not the standardized variables). Multiplying by 2.3 converts to fractional errors for $\log(T_{\text{eff}})$.

model	evaluation sample	A_V			$\log(T_{\text{eff}})$			$\log g$			[Fe/H]		
		ϵ_{sys}	ϵ_{mae}	ϵ_{rms}	ϵ_{sys}	ϵ_{mae}	ϵ_{rms}	ϵ_{sys}	ϵ_{mae}	ϵ_{rms}	ϵ_{sys}	ϵ_{mae}	ϵ_{rms}
TG	F G=15				<	0.0010	0.0014	<	0.065	0.093			
TG	F G=18.5				<	0.0057	0.0080	<	0.35	0.51			
TG	F G=20				0.0045	0.019	0.027	<	1.14	1.53			
TG-allmet	F G=15				$7.4e^{-4}$	0.0058	0.0078	<	0.53	0.90			
TG-allmet	F G=18.5				$8.2e^{-4}$	0.0083	0.011	<	0.66	0.99			
TG-allmet	F G=20				0.0025	0.019	0.027	-0.14	1.19	1.59			
TM-dwarfs	L G=15				<	0.0017	0.0020				<	0.14	0.24
TM-dwarfs	L G=18.5				<	0.0024	0.0033				-0.037	0.26	0.42
TM-dwarfs	L G=20				<	0.0070	0.0090				<	0.82	1.14
TM-giants	L G=15				<	0.0028	0.0037				<	0.22	0.34
TM-giants	L G=18.5				$3.4e^{-4}$	0.0035	0.0045				<	0.31	0.50
TM-giants	L G=20				0.0011	0.0073	0.0092				-0.11	0.74	1.08
TM-allgrav	L G=18.5				<	0.0052	0.0070				<	0.40	0.60
TM-allgrav	F G=18.5				0.0031	0.012	0.017						
TAG	F G=15	-0.0067	0.072	0.15	0.0014	0.013	0.026	-0.052	0.29	0.57			
TAG	F G=18.5	0.039	0.30	0.45	0.017	0.061	0.094	-0.22	1.10	1.53			
TAM	L G=15	0.032	0.18	0.32	0.0022	0.018	0.029				<	0.46	0.79
TAM	L G=18.5	0.047	0.52	0.74	0.010	0.056	0.080				-0.38	1.34	1.81
TGM	L G=15				$8.3e^{-5}$	0.0007	0.0012	-0.022	0.15	0.26	0.014	0.084	0.19
TGM	L G=18.5				$8.8e^{-4}$	0.0043	0.0060	-0.11	0.84	1.15	-0.045	0.35	0.60
TGM	F G=18.5				$2.3e^{-3}$	0.0089	0.014	-0.06	0.59	0.86			




Figure 11. Predictions of the full forward model for the TG problem as a function of $\log(T_{\text{eff}})$ with $\log g$ fixed at 4.0 dex for 12 different bands (the central wavelength of which is indicated at the top of each panel in nm). The black crosses are the (noise-free) grid points, the small red stars are the forward model predictions (at randomly selected AP values) and the blue circles the noisy ($G=15$) grid points. The photon counts plotted on the ordinate are in standardized units.

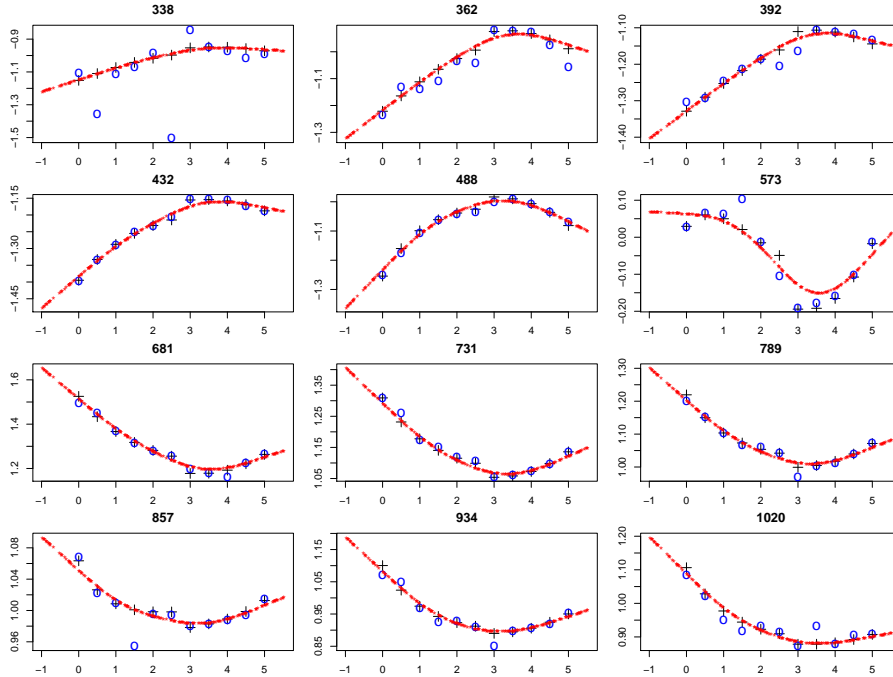


Figure 12. As Fig. 11, but now showing predictions of the full forward model as a function of $\log g$ at constant $T_{\text{eff}}=5000$ K.

It is hard to fairly compare the performance with nearest neighbours (NN), because on the one hand NN is limited by the density of its template grid, but on the other hand it can report the *exact* AP values. If we used a full, noise-free template grid and noisy evaluation objects, then provided the noise is low enough, 1-NN will give exact results. If we instead split the template and evaluation sets to have no common objects, then the precision of at least one of the APs is limited by the grid spacing. We could instead average over the k nearest neighbours, but then NN does badly because it averages over a wide range of the weak AP (a shortcoming which partly motivates ILIUM). To give some comparison, however, I estimate the parameters of all 274 stars in the TG grid of noise-free spectra via leave-one-out cross validation. The errors in both APs are six times larger than the ILIUM result at $G=15$. This at least confirms that ILIUM overcomes the grid resolution limitation of NN.

The above results are for stars with $[\text{Fe}/\text{H}]=0$ dex. I retrained and evaluated ILIUM on a grid with the full range of metallicities (-4 to 1.0 dex; TG-allmet). The errors averaged over this full data set at $G=15$, reported in line 4 of Table 2, are 6–8 times larger than the solar metallicity case. This is entirely due to the new metallicity variance which is not accounted for (modelled) by ILIUM and so is a confusing factor. (A T_{eff} accuracy of 1% at $G=15$ and 2% at $G=18.5$ is nonetheless good for stars which a priori show variance over the full range of T_{eff} , $\log g$ and $[\text{Fe}/\text{H}]$). Curiously, if we apply the model to $G=20$ data, then we see that the performance is no worse than when we limited the problem to solar metallicity stars (compare lines 3 and 6 of Table 2). This is because the photometric noise dominates over the variance introduced by the (unmodelled) metallicity range.

If the AP residuals had a Gaussian distribution, then $\epsilon_{\text{rms}} = 1.25\epsilon_{\text{mae}}$ in Table 2. The fact that ϵ_{rms} is always larger

(sometimes much larger), indicates that there are outliers, which justifies the use of ϵ_{mae} as a more robust error statistic.

In terms of overall error, an SVM does better than ILIUM on the TG-allmet problem at all three magnitudes (Table B1). It's not clear why this is so, given that ILIUM was better on the problem limited to solar metallicity (see section B).

5 APPLICATION TO THE $T_{\text{eff}}+[\text{Fe}/\text{H}]$ PROBLEM (TM)

I now use ILIUM to estimate T_{eff} and $[\text{Fe}/\text{H}]$ on the two grids TM-dwarfs and TM-giants defined in section 3.2. In each case we effectively assume we already have a rough $\log g$ estimate. The remaining spread in $\log g$ in each case acts as cosmic scatter. The models are then applied to the corresponding evaluation sets with noise levels at three magnitudes. The summary performance statistics are shown in six lines in Table 2. As there is essentially no sensitivity to metallicity at T_{eff} above 7000 K (Fig. 9), I only report results on cooler stars, even though the ILIUM models were fit to the full T_{eff} range.

Looking first at the T_{eff} performance, we see that the precision is twice as good as TG at all magnitudes (e.g. 0.007 dex compared to 0.019 dex). This just indicates that we can estimate T_{eff} more precisely for cool stars. Within the T_{eff} range 4000–7000 K there is no strong dependence of the T_{eff} precision with T_{eff} or $[\text{Fe}/\text{H}]$: ILIUM can estimate T_{eff} equally well at all metallicities.

Turning now to metallicity, we see good performance at $G=15$ and $G=18.5$ for dwarfs and giants: random errors of 0.3 dex or less and negligible systematics. At $G=20$ the performance is quite a lot worse (0.7–0.8 dex). There is little dependence of $[\text{Fe}/\text{H}]$ precision or accuracy with $[\text{Fe}/\text{H}]$, as can be seen in the lower panel of Fig. 15: Even for the most metal poor stars in the sample at $[\text{Fe}/\text{H}]=-4.0$ dex the precision is still 0.5 dex. This plot also shows that the AP estimates hardly ever exceed the limits of the training

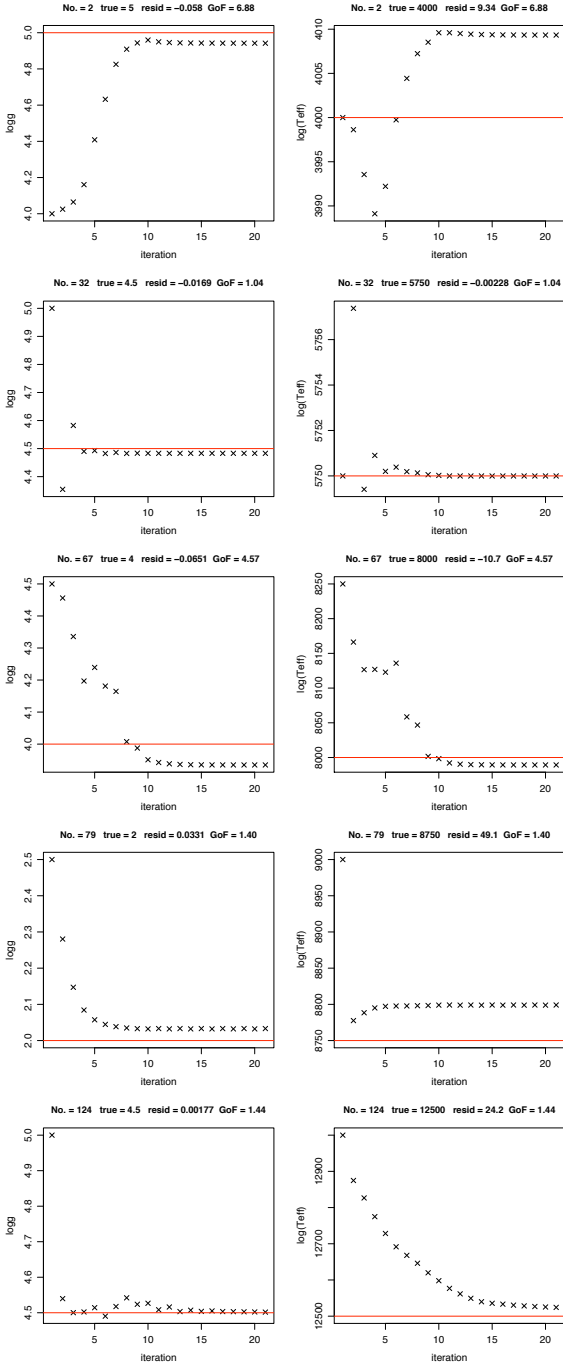


Figure 13. AP evolution for 5 stars in the evaluation data set at $G=15$ ($\log g$ left, T_{eff} right) for the TG problem. The true APs are written at the top of each panel pair and plotted as the red horizontal line. GoF is the reduced χ^2 goodness-of-fit (equation 12).

grid, even though they are allowed to (section 2.5.7), again suggesting natural convergence properties of the algorithm.

We gain some insight into how ILIUM works if we include hot stars in the evaluation set. There is now a strong systematic error in the $[\text{Fe}/\text{H}]$ estimates, as can be seen in the upper panel of Fig. 15. The reason is that the sensitivity of all the spectral bands to metallicity is essentially zero in hot stars. In that case there is no contribution from metallicity to the flux updates in the ILIUM algorithm,

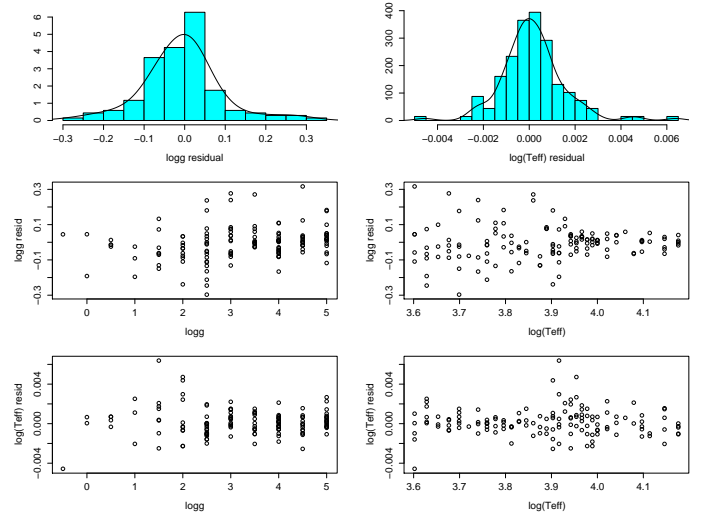


Figure 14. AP residuals for the TG problem at $G=15$




Figure 15. $[\text{Fe}/\text{H}]$ residuals for the TM-dwarfs model at $G=15$ for the full T_{eff} range (top) and for cool stars only ($T_{\text{eff}} \leq 7000 \text{ K}$; bottom). Note the different scales on the ordinate. The systematic error (due to the lack of sensitivity to $[\text{Fe}/\text{H}]$ for hot star spectra) vanishes in the lower panel.

so the flux prediction is entirely from the strong component of the forward model. That predicts a flux corresponding to the average value of the metallicity (Fig. 4) in the training grid, so this is the AP value which ILIUM reports. This is obviously higher than the lowest metallicities, with the result that ILIUM overestimates $[\text{Fe}/\text{H}]$. This is logical and desirable: ILIUM reports the average value in the training data when the spectrum provides no information. The AP distribution is acting as a prior.

The above results have implicitly assumed the $\log g$ of the star to be known well enough to identify it as a dwarf or giant. We saw from the results in the previous section that ILIUM itself can do this (even if the metallicity is not known: line 4 of Table 2). But what if we relaxed this assumption? To test this, I trained and evaluated a new model, TM-allgrav, on the full range of $\log g$ (see section 3.2). Even at $G=18.5$ this model can estimate $[\text{Fe}/\text{H}]$ to a precision of 0.4 dex, which would be enough to trace the metallicity distribution in the Galaxy and identify very metal-poor stars.

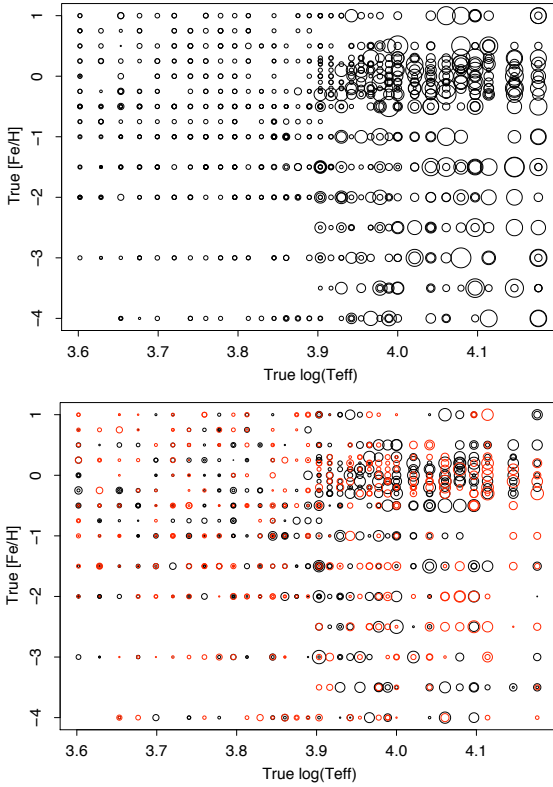


Figure 16. Predicted AP uncertainties (top) and actual residuals (bottom) for $\log(T_{\text{eff}})$ from the TM-dwarf problem at $G=20$. The area of the plotted circle (not the diameter) is proportional to the size of the uncertainty/residual. In the lower panel, positive residuals are shown in black and negative in red. Points below some size limit are not plotted at all.

Section 2.7 gave a formula for the covariance of the APs estimated by ILIUM. For this TM problem, \mathbf{C}_ϕ is a 2×2 matrix with elements $c_{jj'}$. The corresponding uncertainty estimates, $\sqrt{c_{jj'}}$, are plotted as a function of the two APs in the upper panel of Fig. 16. This can be compared to the actual errors (residuals) in the lower panel. The fact that they broadly agree indicates that the uncertainty estimation is useful. This is important, because for many purposes knowing the uncertainty in an estimate is as important as the estimate itself.

6 APPLICATION TO 3-AP PROBLEMS (TAG & TAM)

6.1 Extension to higher dimensions

So far the forward model has comprised two 1D components. For more than two APs the forward modelling approach described in section 2.4 generalizes almost trivially. The principle is to retain the partition of the APs into the two categories “strong” and “weak”. With N_s strong APs, the strong component of the forward model is a single N_s -dimensional function, fit at the n_s unique combinations of the strong APs in the training grid. At any one of these points, k , there are $n_w(k)$ grid points which vary over the N_w weak APs. The mean flux over these is used to fit the strong component. The weak component at point k is then built by making an N_w -dimensional fit to the flux residuals (with respect to the strong component) over the $n_w(k)$ points at k . Thus we have n_s independent weak compo-

nents. The components are combined and applied as in the 1D+1D case. Nothing else in the ILIUM algorithm is changed.

Here I apply ILIUM to two different problems both with two strong APs and one weak AP. This is especially important for stellar parametrization (and Gaia) because in practice we have to accommodate variable interstellar extinction, which can be large at low Galactic latitudes and has at least a large an impact as T_{eff} on the spectrum (see Fig. 7).

I fit the 2D strong component of the forward model using a thin plate spline, a type of smoothing spline closely related to kriging (e.g. Hastie et al. 2001). The number of degrees of freedom is set to $n_s/2 = 330/2 = 155$ in both problems. The 1D weak components are again fit using smoothing splines with the dof set as before (section 2.5.2).

6.2 Results for $(T_{\text{eff}}+A_V)+\log g$ (TAG)

In this problem the forward model is fit over the full range of T_{eff} , A_V and $\log g$ for solar metallicity stars (TAG in section 3.2). The 1D cut through the fit in Fig. 17 shows that the newly introduced A_V variance is fit accurately (as are the variations in the other APs). The summary performance when applying ILIUM with this model to $G=15$ spectra is shown near the bottom of Table 2. Compared to the results on the TG data set at the same magnitude, the errors in $\log(T_{\text{eff}})$ and $\log g$ are, at 0.013 dex (3%) and 0.3 dex respectively, considerably worse. This is due to the extra variance introduced by the very wide range of extinction. Yet these errors are still small enough to be scientifically useful, especially when we note that the systematics are quite small. The new AP, extinction, can be estimated very well: a mean absolute error of just 0.07 mag. If we limit the analysis (without changing the fitted model) to just low extinction stars ($A_V \leq 1.0$ mag), then A_V can be estimated only slightly better (0.056 mag) but $\log(T_{\text{eff}})$ more so (0.008 dex).

At $G=18.5$, the errors of course increase (Table 2): while T_{eff} and A_V are still manageable, at 1.1 dex $\log g$ is not. Fortunately Gaia will provide parallaxes for many stars to improve the $\log g$ estimates (given the T_{eff} and A_V estimates from ILIUM).

At $G=15$ the SVM has comparable performance on this problem, but is somewhat better at $G=18.5$ (see section B).

6.3 Results for $(T_{\text{eff}}+A_V)+[\text{Fe}/\text{H}]$ (TAM)

We now swap $\log g$ for $[\text{Fe}/\text{H}]$ and train and evaluate ILIUM on the TAM problem (which is a dwarf sample: the results are broadly similar for a giant sample). The summary performance at $G=15$ is listed near the bottom of Table 2. The performance is degraded significantly compared to the case with no A_V variance (TM-dwarfs), much more so for T_{eff} than for $[\text{Fe}/\text{H}]$. At 0.46 dex the metallicity error is acceptable. Alas at $G=18.5$ this degrades to an almost useless level with the additional problem of a large systematic.⁶

If we limit the analysis (at $G=15$) to just low extinction stars ($A_V \leq 1.0$ mag), then the $[\text{Fe}/\text{H}]$ precision is essentially unchanged (0.44 dex). This implies that it is no more difficult (on average) to determine the metallicity of stars with high extinction than of stars with low extinction. So why is the error larger here than in the zero extinction case reported for the TM problem? It is because we now

⁶ Contrary to expectations, systematic trends can rarely be corrected for. We would have to plot the residual vs. the *estimated* AP, and if the scatter is larger than the systematic trend then a correction cannot be made.

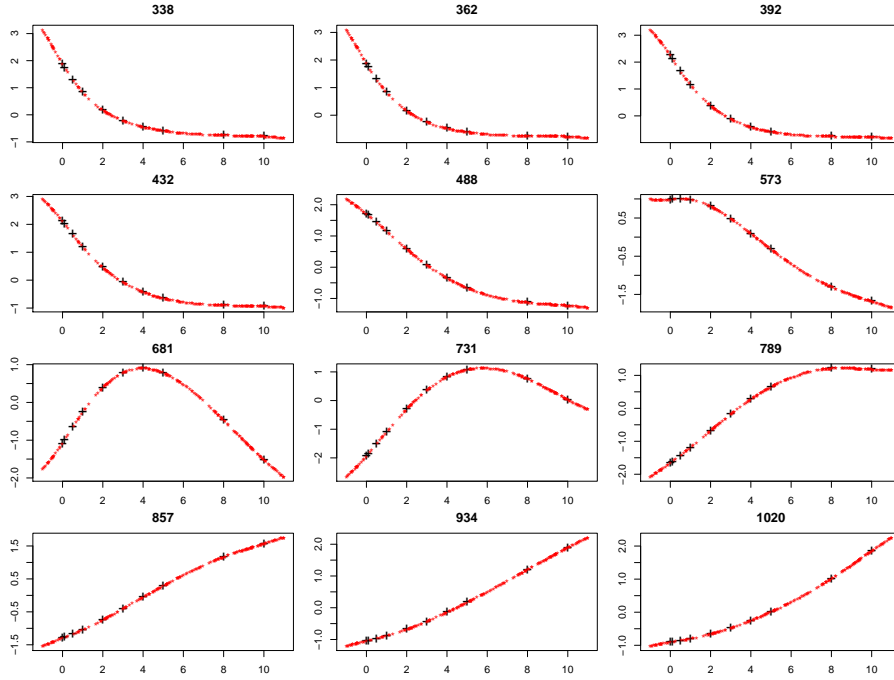


Figure 17. Predictions (small red stars) of the full forward model for the TAG problem as a function of A_V with T_{eff} fixed at 10000 K and $\log g$ at 4.0 dex for 12 different bands (the central wavelength of which is indicated at the top of each panel in nm). The black crosses are the (noise-free) grid points. The flux plotted on the ordinate is in standardized units.

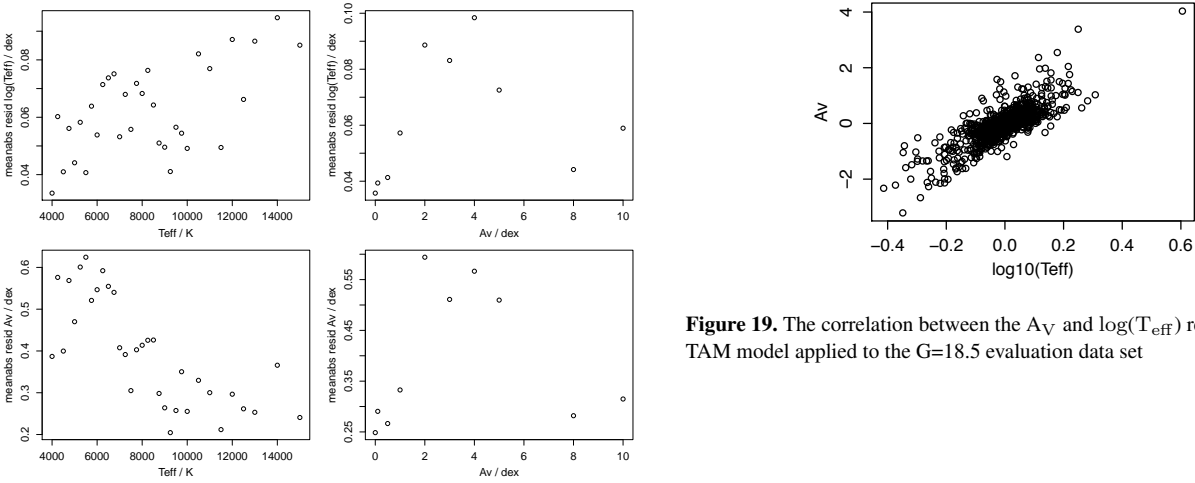


Figure 18. The trend in the mean absolute errors for $\log(T_{\text{eff}})$ and A_V in the TAM problem on the G=18.5 evaluation data set (which spans the full range of T_{eff} , A_V and $[\text{Fe}/\text{H}]$)

have a large extinction range *a priori*, so an uncertainty in determining A_V corresponds to an uncertainty in $[\text{Fe}/\text{H}]$. As $[\text{Fe}/\text{H}]$ is weak, a relatively small uncertainty in A_V is magnified into a larger one in $[\text{Fe}/\text{H}]$.

Note that the extinction errors are larger than those reported with the TAG problem (0.18 mag compared to 0.07 mag at G=15). This is a consequence of having limited the analysis to cool stars. Indeed, the mean absolute error on the present problem is only 0.1 mag if we analyse the hot stars ($T_{\text{eff}} > 7000$ K) rather than the

Figure 19. The correlation between the A_V and $\log(T_{\text{eff}})$ residuals for the TAM model applied to the G=18.5 evaluation data set

cool ones: A_V can be estimated more accurately for hotter stars. This is not surprising because hot star spectra are simpler (e.g. no metallicity signature) so it should be easier to untangle the effects of temperature and extinction. This is illustrated in Fig. 18 at G=18.5, which plots the dependence of the $\log(T_{\text{eff}})$ and A_V residuals with these parameters (averaged over all $[\text{Fe}/\text{H}]$). It also shows that T_{eff} is estimated more precisely for cool rather than hot stars, as found also for the TM problem.

Nonetheless, the T_{eff} and A_V errors are quite large at G=18.5: 8–20% for T_{eff} and 0.2–0.6 mag in A_V (Fig. 18). Moreover, their residuals show a strong positive correlation (Fig. 19): a tendency to overestimate one results in an overestimation of the other. This suggests that these two APs are degenerate in these low resolution BP/RP spectra, something which Fig. 7 also suggests. Let us now investigate this further.

6.4 Identification of a $T_{\text{eff}}-A_V$ degeneracy

Like most estimation algorithms, ILIUM simply tries to find the best single solution (plus an uncertainty estimate). If we consider a likelihood function of the data given the APs, $P(\text{Data}|\phi)$, then the algorithm is trying to find the maximum of this, plus some measure of the width of this peak. But if there is a degeneracy in the APs then the peak position and width are an inadequate summary of the likelihood. A degeneracy can be defined as two or more stars which have spectra differing by an amount consistent with the photometric noise. On observing one of these spectra we would be unable to distinguish between the AP solutions. This applies trivially to stars with very small AP differences (thus the need to quote an uncertainty). What we are interested in here are objects which have significantly different APs (a large fraction of the grid range), such that a simple approximation of the likelihood function (e.g. a Gaussian in the APs with covariance given by equation 11) fails to represent the uncertainties accurately.

As ILIUM is a local search method, it is a priori possible that when initialized at different points it could converge on different (degenerate) solutions. To explore this I ran ILIUM 25 times for each star with the initialization chosen at random each time, rather than by the nearest neighbour. (I used the TAG model with $G=18.5$ spectra.) Fig. 20 plots the various solutions in the $T_{\text{eff}}-A_V$ plane for three different stars, chosen to illustrate degeneracy.

For the first star (top panel), we see that 22 of 25 initializations over quite a range of T_{eff} and A_V converge on a small region (slightly offset from the true solution possibly due to noise in the spectrum). The three “wrong” solutions occur because the algorithm is initialized too far from a good solution: the searching method gets stuck in a poor local minimum. Inspection of their predicted spectra shows them to be very different from the true spectrum, with GoF values (equation 12) of more than a few hundred. Such a high GoF would be used in practice to flag and reject such solutions.

Turning to the second star (middle panel) we see two distinct regions of convergence, one of which is close to the true APs. Are the other solutions at around $T_{\text{eff}} = 13\,000$ K simply poor (bad convergence)? When we inspect the predicted spectra (not shown), we see this is not the case. The two sets of spectra are indistinguishable with the noise, all having similar and small GoF values (0.01–15).

This is better illustrated by the third star (bottom panel), where we see a lack of convergence on one or two solutions in favour of a complete ridge of solutions. Fig. 21 plots the initial and final (predicted) spectra for these 25 runs of ILIUM. All but one or two of these final spectra agree very closely with the true spectrum, even though they have very different APs. So this is a true degeneracy, and not poor convergence of ILIUM. This degeneracy is also reflected by a high value for the expected correlation coefficient between T_{eff} and A_V calculated by equation 11. Identification of similar ridges in AP space of near-identical spectra for other stars suggests that the $T_{\text{eff}}-A_V$ degeneracy is widespread.

6.5 Mapping of the $T_{\text{eff}}-A_V$ degeneracy

We can map this degeneracy systematically using just the forward model to generate spectra on a fine grid of T_{eff} and A_V . For each predicted spectrum I adopt as its sigma spectrum that of the star in the original grid which has the closest APs. The expected Mahalanobis distance between a hypothetically measured spectrum, \mathbf{p} ,




Figure 20. Multiple random initialization of ILIUM for three different stars (in the three panels). In each panel the starting points are shown as green dots connected to the solutions (after 20 iterations) shown as blue triangles. The true APs are shown as a red cross.

and any noise-free spectrum, \mathbf{p}' , in this grid is, with $\delta\mathbf{p} = \mathbf{p} - \mathbf{p}'$,

$$D^2 = \delta\mathbf{p}^T \mathbf{C}_p^{-1} \delta\mathbf{p} \quad (13)$$

$$= \sum_{i=1}^{i=I} \left(\frac{p_i - p'_i}{\sigma_{p_i}} \right)^2 \quad (14)$$

the simplification following because in these simulations there is no inter-pixel noise correlation. A degeneracy arises between two stars when D^2 is sufficiently small that it could arise just from photometric noise. Under the null hypothesis (H_0) that the differences between the spectra are only due to Gaussian noise and that each pixel is independent, D^2 follows a χ^2 distribution with $I - 1 = 67$ degrees of freedom. I will define two stars as non-degenerate only if the probability of observing their given distance or more under H_0 is 1% or less. This corresponds to a critical value of $D_{\text{lim}}^2 = 96.8$

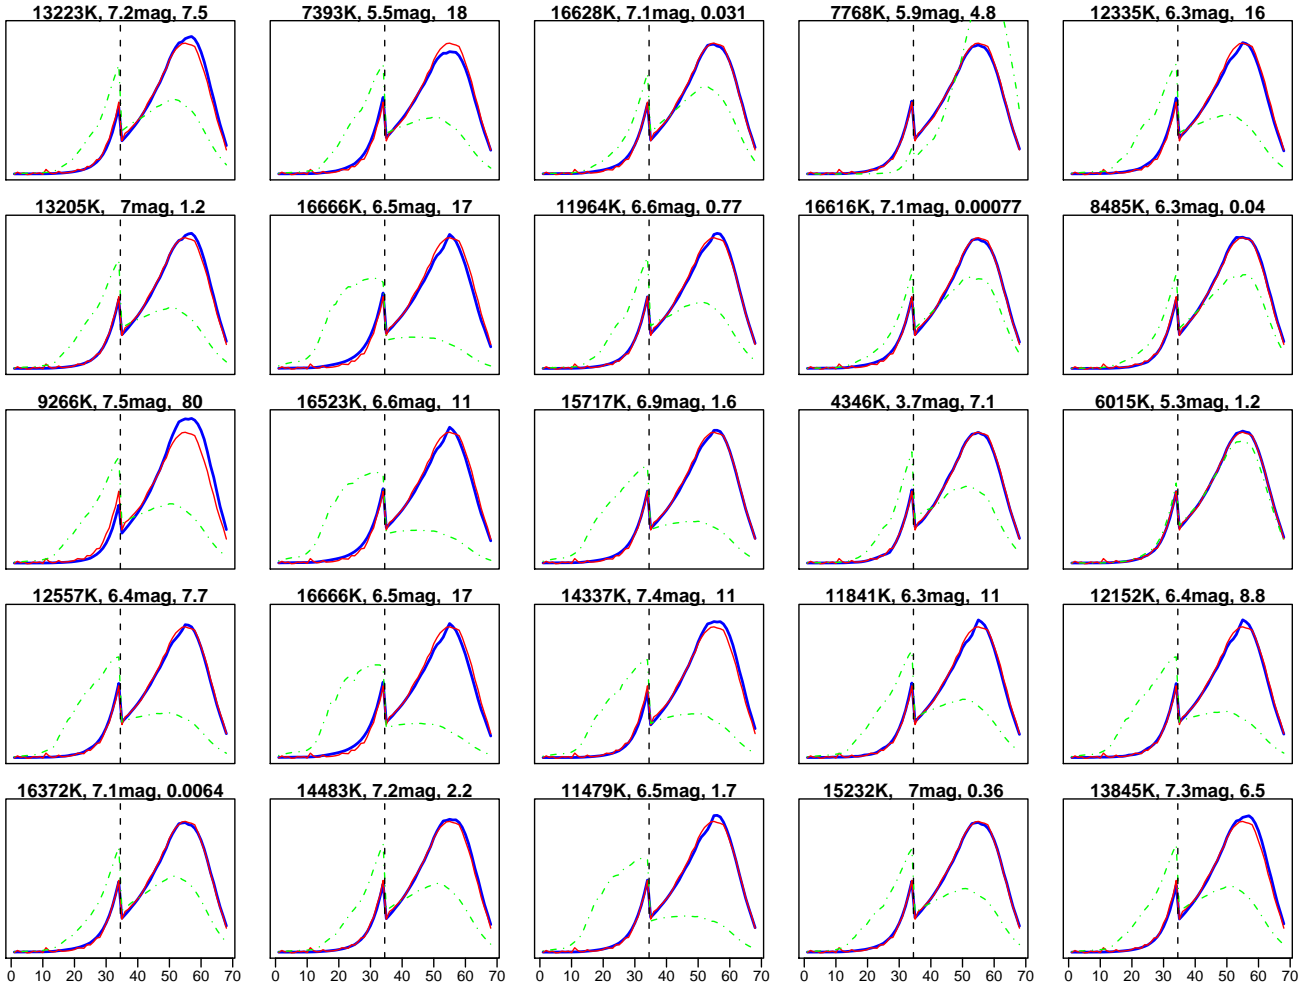


Figure 21. The initial (green dashed-dot line), final (blue thick line) and true (red thin line) spectra for the 25 runs of ILIUM for the one star shown in the bottom panel of Fig. 20, plotted vs. band number (BP and RP are separated by the vertical dashed line). The T_{eff} and A_V of the ILIUM solutions (the final spectra) are plotted at the top of each panel along with the GoF of each. The true APs are 5500 K and $A_V = 5.0$ dex

(or a reduced- χ^2 value of 1.44). In other words, stars separated by a smaller distance are considered degenerate.⁷

For each star in the fine grid I calculate the distance to all of the other stars. The corresponding probabilities (of getting that distance or more) can then be plotted as a *degeneracy map* in the $T_{\text{eff}}-A_V$ plane. Maps for 15 stars are shown in Fig. 22 (similar patterns are observed throughout the grid). We see that the APs are correlated, the high probability region extends over a large range of the grid and in some cases is multimodal. In other words, there is a strong degeneracy over the whole plane. What this means in practice is that if we observed a (noisy) spectrum, p , at the position of one of the crosses in Fig. 22, then this spectrum is indistinguishable from all the spectra lying within the contours, on account of the noise.

⁷ χ^2 is the distribution followed by a sum of squares of independent unit Gaussian variables, $N(0, 1)$, and has a (unnormalized) density function $P'(D^2) = D^{\nu-2} e^{-D^2/2}$ where ν is the degrees-of-freedom. $\nu = I - 1$ because one degree-of-freedom is “lost” by the fact that they have a common G magnitude, although this is of no practical significance here. D_{lim}^2 is defined by $P(D^2 \geq D_{\text{lim}}^2 | H_0) = \int_{D_{\text{lim}}^2}^{\infty} P' dD^2 = 0.01$. Fig. 22 plots this probability (and several others) as a contour.

Therefore we also cannot distinguish between the corresponding APs.

We can compare the degeneracy maps with the expected covariance for individual objects calculated from equation 11 by treating the latter as the covariance in a Gaussian likelihood function in the APs. While the estimated correlations agree qualitatively with what we observe in Fig. 22, this Gaussian likelihood approximation does not accurately reproduce the shape of the degeneracy ridges. Indeed it cannot, because the degeneracy ridges are not symmetric about the true estimate, and A_V cannot be negative. This lack of detailed agreement is not very surprising, however, because equation 11 is a valid approximation only when the linear relation in equation 1 holds, i.e. when the errors are small.

The implication of this is significant: It is misleading to report a single estimate for T_{eff} and A_V (even if accompanied by a simple covariance estimate). Rather, we must report a whole ridge of solutions. As we can map these degeneracies in advance, the purpose of a classifier such as ILIUM would be to identify *a* solution, which we use to identify the corresponding degeneracy ridge (either as a probability grid or a fitted approximation to it). This study shows that a single nearest neighbour initialization of ILIUM is usu-




Figure 22. Degeneracy map (a likelihood function) for the forward model of the TAG grid at $G=18.5$. For each of the 15 stars (panels) marked as a red cross, each contour indicates those stars which have a probability P of having a spectrum indistinguishable to within the expected photometric noise. Contours are plotted for $\log P = \{-4, -3, -2, -1, 0\}$ (due to the steep sides of the χ^2 function, these are barely distinguishable).

ally adequate to allow ILIUM to find a good solution and therefore the correct ridge. The minority of cases where no good solution is found are usually identifiable by the large GoF value. In those cases we may want to re-initialize ILIUM from some other point or improve the convergence mechanism.

Note that Fig. 22 is for $G=18.5$. At $G=20$ the degeneracy regions are somewhat larger. At $G=15$ the SNR is sufficiently high that there is no longer any degeneracy on a scale of practical relevance: the contours in the degeneracy map are smaller than the grid point separation (0.2 mag in A_V and 0.01 dex in $\log(T_{\text{eff}})$). This is consistent with the small errors in these APs recorded in Table 2.

6.6 Including additional or prior information

Additional information could help to reduce the extent of the degeneracy.

The parallax reported by Gaia combined with the apparent G -band magnitude gives an estimate of $M_G + A_G$, where M_G is the absolute magnitude in the G band and A_G is the extinction in the G -band (which we could estimate directly from ILIUM instead of A_V). Combined with an HR-diagram (which is just a prior probability density function over T_{eff} and M_G) this provides additional constraints on T_{eff} and A_G . A method for doing this is outlined in Bailer-Jones (2010).

We could set weak prior probabilities on the APs – the probability of the APs unconditional on the BP/RP spectrum, parallax or G -band magnitude – based on a simple model of the Galaxy. For example, we know that cool stars are much more common than hot

stars, so an unreddened cool star is a priori more likely than a highly extinct OB star. Likewise, if the star is observed at high Galactic latitude we could confidently assign a lower prior probability to high extinctions. This latter prior is particularly useful, because we saw in sections 4 and 5 that if the extinction can be fixed then ILIUM gives much better estimates. Normally, however, we wouldn't want to bias the Gaia AP estimates with *detailed, current* knowledge of Galactic structure, given that the primary goal of Gaia is to improve this knowledge.

For very bright stars ($G \lesssim 12$; Gaia will saturate at around $G=6$) we may have AP estimates coming from the high-resolution RVS spectrograph on Gaia measuring the CaII triplet at 860 nm (Katz et al. 2004). In principle we can estimate T_{eff} independently of A_V using these data. Note that because the ridges in Fig. 22 have a low inclination, additional information on T_{eff} is more useful than information on A_V .

Finally, estimates of APs from other surveys could be incorporated and combined with the Gaia ones, provided they can be interpreted as probability distributions.

Given the opportunity, one should design an ad hoc spectrophotometric system to address the specific goals of the survey. Particularly desirable is a system which maximizes the separability of the APs (a method for doing this was presented by Bailer-Jones 2004). This opportunity was taken by the Gaia community, which designed a multiband photometric system to address the specific Gaia scientific goals (Jordi et al. 2006). But this was later replaced with the present (and less desirable) BP/RP spectrophotometry, ultimately for financial reasons.

7 FURTHER APPLICATIONS (AND THE TGM PROBLEM)

In the previous section the forward model was extended to be 2D in the strong APs and 1D in the weak. In a similar way it can be adapted to model one strong AP (e.g. T_{eff}) and two weak APs (e.g. $[\text{Fe}/\text{H}]$ and $\log g$) simultaneously. This is of practical use in situations where extinction can be assumed to be low. I used this to fit a forward model to the TGM grid (section 3.2) and applied ILIUM as before. The results are shown in the bottom section of Table 2. At $G=15$, the performance in T_{eff} and $[\text{Fe}/\text{H}]$ is considerably better than obtained with the corresponding TM problems in which there was unmodelled $\log g$ variation: we now obtain 0.3% in T_{eff} for the full T_{eff} range and 0.1 dex in $[\text{Fe}/\text{H}]$ for cool stars. $\log g$ is also better than found with TG-allmet, but not with TG where metallicity was fixed to zero. This as we would expect: at high SNR, modelling the extra variance improves the results in all APs.

This is not the case at $G=18.5$. Here the performance of the TGM model on all three parameters is similar to that obtained in the TG-allmet or TG-allgrav problems on the appropriate T_{eff} ranges. I suspect that the extra variance introduced by the noise is complicating the problem, so that the modelling the additional AP (rather than marginalizing over it) brings no improvement in performance (but the obvious advantage that we now determine all three APs instead of just two.) It is also worth pointing out that when using the degeneracy mapping method described in section 6.5, I find there is a very strong and complex degeneracy between $[\text{Fe}/\text{H}]$ and $\log g$ at $G=18.5$.

ILIUM could be further extended to more or other APs or to different problems, provided the partition between strong and weak APs can be retained. If such a clear distinction were not possible then a different approach to forward modelling would be necessary. If we had more than two or three weak or strong APs, then the smoothing splines are unlikely to provide adequate fits, unless we had a lot of data. To overcome this we would need more structured regression models suited for fitting high dimensional sparse data sets (e.g. neural networks).

8 SUMMARY AND CONCLUSIONS

I have introduced an algorithm for estimating parameters from multi-dimensional data. It uses a forward model of the data to effectively perform a nonlinear interpolation of a template grid via the Newton-Raphson method. It is intended to overcome both the non-uniqueness issue of direct modelling of inverse problems as well as the finite grid density and metric definition issues of k nearest neighbours. ILIUM makes use of the sensitivity of the data to the astrophysical parameters to find an optimal solution. This is convenient, because in principle it means we don't need to worry too much about feature selection: low sensitivity features will automatically be down weighted. An important component of the algorithm is the division of the forward model into two parts which independently model the variance of the "strong" APs (such as T_{eff} and A_V) and the weak APs (such as $\log g$ and $[\text{Fe}/\text{H}]$). Good fits could be obtained using low-dimensional (1 or 2) smoothing splines. As it is based on a forward model, ILIUM naturally provides AP uncertainty estimates (actually full covariances) and a goodness-of-fit, in contrast to most inverse modelling methods.

I applied ILIUM to the problem of estimating APs from simulations of the low resolution Gaia spectrophotometry, BP/RP. Results are summarized in Table 2. When limited to zero extinction

stars, T_{eff} can be estimated at $G=15$ to a (mean absolute) accuracy of better than 1% when the metallicity and surface gravity are entirely unknown ($\log g = -0.5$ to 5.0 dex; $[\text{Fe}/\text{H}] = -4$ to $+1.0$ dex). At $G=18.5$ and $G=20.0$ (the limiting Gaia magnitude) the average T_{eff} accuracy over the full range of APs is 2% and 4% respectively. At $G=15$ $\log g$ can be estimated to 0.5 dex if $[\text{Fe}/\text{H}]$ is entirely unknown but to 0.15 dex if $[\text{Fe}/\text{H}]$ is also modelled (full range of the three APs). Limiting to solar metallicity stars this improves to better than 0.1 dex and is still 0.35 dex at $G=18.5$. $[\text{Fe}/\text{H}]$ for cool stars can be estimated to 0.1 dex at $G=15$ and to 0.15–0.35 dex at $G=18.5$, the better result obtained if we know the star is a dwarf. At $G=20$ $\log g$ and $[\text{Fe}/\text{H}]$ cannot be estimated to any useful accuracy. Of course, for population studies we can average over many stars to reduce the population estimate of metallicity, limited by the systematic errors and any correlations in the data.

If we extend the strong component of the forward model to simultaneously model A_V over a very wide range (0–10 mag), then the performance on the weak APs is significantly degraded on account of this extra variance, such that reasonable accuracies can be reached at $G=15$ but not at $G=18.5$. A_V itself can be estimated remarkably well, however, 0.07–0.2 mag at $G=15$ and 0.3–0.5 mag at $G=18.5$. The extra variance also affects the T_{eff} determination. However, these statistical errors do not tell the whole story: I have shown that there is a strong and ubiquitous degeneracy between T_{eff} and A_V intrinsic to the BP/RP spectra. Thus in addition to reporting single optimal estimates, the Gaia catalogue will need to provide the corresponding degeneracy map, which can be built in advance using the forward model. The degeneracies could be reduced (and the weak APs also then estimated more accurately) if we can use additional information. Possible sources include the apparent magnitude and parallax measured by Gaia, external data, and/or weak priors from a very simple Galaxy and stellar evolution model.

For some problems ILIUM outperformed a support vector machine in terms of smaller residuals, but in other cases the SVM was better. This should be assessed in more detail after improvements to the ILIUM algorithm have been explored. For example, the updating method is very simple – stopping after a fixed number of iterations – whereas a more adaptive approach may help on larger variance problems (e.g. noisier data). We may also want the AP update weighting to take into account the noise (and not just the sensitivity).

Now that we have a forward model, further possibilities for modelling open up. ILIUM is just a method for locating the best APs. If we define a distance metric (such as equation 13), then we can define a likelihood function, $P(D|\phi)$, which when combined with a suitable prior defines a (non-analytic) posterior over the APs, $P(\phi|D)$. We can then use one of many sampling methods, such as Markov Chain Monte Carlo, to sample this as a function of ϕ , thus yielding a complete Bayesian probabilistic solution which varies smoothly over the APs. This approach to parameter estimation has been used in several areas of astronomy, such as galaxy classification (e.g. Heavens et al. 2000) and inference of cosmological parameters (e.g. Percival et al. 2007). (If we dispensed with the forward model and just calculated probabilities at points in the original grid, then the principle is similar to that used by Shkedy et al. 2007.) While offering some advantages, this approach is much slower than ILIUM, because for the present application it would require of order 10^4 to 10^5 samples and hence this many evaluations of the forward model, compared to about 10^2 for ILIUM.

ACKNOWLEDGEMENTS

I would like to thank Anthony Brown, Lennart Lindegren and Carola Tiede for useful discussions. This work makes use of Gaia simulated observations, which have been produced thanks to the efforts of many people in the Gaia DPAC. In this respect I would particularly like to thank Anthony Brown, Yago Isasi, Xavier Luri, Paola Sartoretti, Rosanna Sordo and Antonella Vallenari, without whose efforts the predictions of the Gaia performance would not have been possible. The GOG simulations were produced using the MareNostrum supercomputer at the Barcelona Supercomputing Center – Centro Nacional de Supercomputación. I am also grateful to the Marcs team at Uppsala University for producing new stellar spectral simulations for Gaia data processing purposes.

REFERENCES

- Bailer-Jones C.A.L., 2004, *A&A* 419, 385
- Bailer-Jones C.A.L., in *The Three-dimensional universe with Gaia*, C. Turon, K.S. O’Flaherty, M.A.C. Perryman (eds), ESA, SP-576, 393
- Bailer-Jones C.A.L., 2009, Gaia DPAC Technical note, GAIA-C8-TN-MPIA-CBJ-042
- Bailer-Jones C.A.L., 2009, Gaia DPAC Technical note, GAIA-C8-TN-MPIA-CBJ-043
- Bailer-Jones C.A.L., 2009, Gaia DPAC Technical note, GAIA-C8-TN-MPIA-CBJ-046
- Bailer-Jones C.A.L., 2009, Gaia DPAC Technical note, GAIA-C8-TN-MPIA-CBJ-048
- Bailer-Jones C.A.L., 2010, Gaia DPAC Technical note, GAIA-C8-TN-MPIA-CBJ-049
- Bailer-Jones C.A.L. (ed.), 2008, *Classification and Discovery in Large Astronomical Surveys*, AIP Conference Proceedings vol. 1082, American Institute of Physics
- Bailer-Jones C.A.L., Smith K.S., Tiede C., Sordo R., Vallenari A., 2008, *MNRAS* 391, 1838
- Ball N.M., Brunner R.J., Myers A.D., Strand N.E., Alberts S.L., Tchong D., 2008, *ApJ* 683, 12
- Bishop C.M., 2006, *Pattern recognition and machine learning*, Springer
- Burges C.J.C., 1998, *Data mining and knowledge discovery* 2, 121
- Brown A.G.A., 2006, Gaia DPAC Technical note, GAIA-CA-TN-LEI-AB-005
- Brown W.R., Allende Prieto C., Beers T.C., Wilhelm R., Geller M.J., Kenyon S.J., Kurtz M.J., 2003, *AJ* 126, 1362
- Cardelli J.A., Clayton G.C., Mathis J.S., 1989, *ApJ* 345, 245
- Cortes C., Vapnik V., 1995, *Machine Learning* 20, 273
- Firth A.E., Lahav O., Somerville R.S., 2003, *MNRAS* 339, 1195
- Gao D., Zhang Y.-X., Zhao Y.-H., 2008, *MNRAS* 386, 1417
- Gustafsson B., Edvardsson B., Eriksson K., Jrgensen U.G., Nordlund A., Plez B., 2008, *A&A* 486, 951
- Hastie T., Tibshirani R., Friedman J., 2001, *The elements of statistical learning*, Springer
- Hawley S.L., Covey K.R., Knapp G.R., Golimowski D.A., Fan X., Anderson S.F., Gunn J.E., Harris H.C., et al., 2002, *AJ* 123, 3409
- Heavens A.F., Jimenez R., Lahav O., 2000, *MNRAS* 317, 965
- Huertas-Company M., Rouan D., Tasca L., Soucail G. Le Fèvre O., 2008, *A&A* 478, 971
- Isasi Y., Zaldua I., Sartoretti P., Luri X., Babusiaux C., Masana E., Gaia DPAC Technical note, 2009, GAIA-C2-UG-UB-YI-003-6
- Jordi C., Hoeg E., Brown A.G.A., Lindegren L., Bailer-Jones C.A.L., Carrasco J.M., Knude J., Straizys V., et al., 2006, *MNRAS* 367, 290
- Katz D., Soubiran C., Cayrel R., Adda M., Cautain, R., 1998, *A&A* 338, 151
- Katz D., Munari U., Cropper M., Zwitter T., Thévenin F., David M., Viala Y., Crifo F., et al., 2004, *MNRAS* 354, 1223
- Lindegren L., Babusiaux C., Bailer-Jones C.A.L., Bastian U., Brown A.G.A., Cropper M., Hog E., Jordi C., et al., 2008, *Proc. IAU Symposia*, vol. 248, 217
- Mignard F., Drimmel R. (eds), 2007, *DPAC Proposal for the Gaia Data Processing*, Gaia DPAC Technical note, GAIA-CD-SP-DPAC-FM-030
- Lejeune Th., Cuisinier F., Buser R., 1997, *A&ASS* 125, 229
- Luri X., Babusiaux C., Masana E., 2005, in *The Three-dimensional universe with Gaia*, C. Turon, K.S. O’Flaherty, M.A.C. Perryman (eds), ESA, SP-576, 357
- O’Mullane W., Lammers U., Bailer-Jones C.A.L., et al., in *Astronomical Data Analysis Software and Systems 16*, R.A. Shaw, F. Hill and D.J. Bell (eds), ASP Conf. Ser. 376, 99
- Percival W.J., Cole S., Eisenstein D.J., Nichol R.C., Peacock J.A., Pope A.C., Szalay A.S., 2007, *MNRAS* 381, 1053
- Recio-Blanco A., Bijaoui A., de Laverny P., 2006, *MNRAS* 370, 141
- Re Fiorentin P., Bailer-Jones C.A.L., Lee Y.S., Beers T.C., Sivarani T., Wilhelm R., Allende Prieto C., Norris J.E., 2007, *A&A* 467, 1373
- Richards G.T., Nichol R.C., A.G. Gray, et al., 2004, *ApJSS* 155, 257
- Shkedy Z., Decin L., Molenberghs G., Aerts C., 2007, *MNRAS* 377, 120
- Snider S., Allende Prieto C., von Hippel T., Beers T.C., Sneden C., Qu Y., Rossi S., 2001, *ApJ* 562, 528
- Salzberg S., Chandar R., Ford H., Murthy S.K., White R., 1995, *PASP* 107, 279
- Sordo R., Vallenari A., 2008, Gaia DPAC Technical note, GAIA-C8-DA-OAPD-RS-002

APPENDIX A: ASSESSING MODEL PERFORMANCE

The ILIUM forward model should obviously be fit on the full available grid and using noise-free data. (We need to achieve a minimum grid density in order to reliably assume that the modelled signal vary smoothly between the grid points, yet the grid need be no denser. Some smoothness assumption must be made by any

method of estimating continuous parameters.) Yet at first sight one may consider it illegitimate to fit ILIUM on the full grid and then assess its performance on (noisy) spectra selected from the same grid. (Note that the initialization set is never included in the evaluation set!) The equivalent approach is generally considered invalid for inverse models, because these are fit by minimizing the AP error on a training set: If they are not properly regularized the model will overfit the training data by learning non-general aspects of these data (such as the noise). But the situation with ILIUM is different, because the best forward model fit is identified (interactively at present) according to its smoothness and the photometry error, not the AP error. That is, the regularization is also done independently of the AP error. It might anyway seem preferable to start with a grid twice as dense as required, fit the forward model on half the data and evaluate ILIUM on the other half. But if the smoothness requirement is met, this will anyway give a similar performance.

The underlying issue here – relevant to all machine learning methods – is how to make a fair assessment of performance. Ultimately all assessments are limited by the fact that training and evaluation data are drawn from a common grid. Yet it would be useless to train the classifier on spectra generated by one set of stellar models and evaluate it on spectra generated by another, because this “performance” would reflect the intrinsic differences between the stellar models. This remains a quandary, also because tests based on synthetic spectra ignore the realities of non-Gaussian noise and cosmic scatter of real data. The most reliable assessment of performance would be to first determine APs of a set of high resolution spectra (using whatever method) based on a particular stellar model and define these APs as true. We would then degrade the spectra to the dispersion, line-spread-function and noise of real Gaia data and compare the ILIUM estimates with the true ones. Such tests are planned but are considerably more arduous than what’s done here.

APPENDIX B: COMPARATIVE PERFORMANCE OF A SUPPORT VECTOR MACHINE

It is not the goal of this article to make a detailed comparison between the accuracy of ILIUM and conventional machine learning methods, but a brief comparison is useful. I therefore applied a support vector machine (SVM) (e.g. Cortes & Vapnik 1995, Burges 1998) to some of the problems presented in the article. The SVM is used to directly model the inverse problem by training on noisy data. Although the SVM training has a unique solution for a given set of data, it has three hyperparameters (the length scale parameter γ and two regularization parameters ϵ and C) which must be optimized (“tuned”), which is just a higher-level training procedure. I optimize these hyperparameters simply via a brute force search of a regular grid over the hyperparameters, typically with 512 or 1000 models (8 or 10 values of each hyperparameter). To achieve this the data set must be split into three independent parts: the training set, a test set (used to select the best combination of hyperparameters) and the evaluation set (on which the final performance is calculated). (If we just use the same data for testing and evaluation – as is frequently done – then the results are somewhat better; unfairly so, because then the SVM is tuned on the very data set on which it is finally evaluated. I nonetheless did this for the TG problem because of the small amount of data: just 274 stars.) For each of the problems/data sets described in section 3.2, the set is randomly split into three, equal-sized disjoint parts to build these subsets. The SVM is tuned on noisy data, separately for each mag-

Table B1. Performance of an SVM on a selection of the problems reported in Table 2. The error statistic is the mean absolute error, ϵ_{mae} . Figures in bold indicate that the error is significantly larger (by more than a third) than the ILIUM error, and vice versa for numbers in italics. Note that the models for the TG problem used just a two-way split into train/test sets due to lack of data, which gives them optimistic results.

model	test sample	A_V	$\log(T_{\text{eff}})$	$\log g$	[Fe/H]
TG	F G=15		0.0015	0.10	
TG	F G=18.5		0.0064	0.47	
TG	F G=20		0.016	0.90	
TG-allmet	F G=15		<i>0.0017</i>	<i>0.091</i>	
TG-allmet	F G=18.5		<i>0.0060</i>	<i>0.41</i>	
TG-allmet	F G=20		<i>0.013</i>	<i>0.77</i>	
TM-dwarfs	L G=15		0.0019		0.20
TM-dwarfs	L G=18.5		0.0064		0.37
TM-dwarfs	L G=20		0.013		0.63
TAG	F G=15	0.084	0.015	0.029	
TAG	F G=18.5	0.24	<i>0.045</i>	<i>0.62</i>	

nitude, and separately for each AP (as an SVM can only model one output).

The SVM results are summarized in Table B1. This obviously conceals details, such as the existence of systematic errors in some cases. For example, T_{eff} is systematically underestimated for $T_{\text{eff}} \geq 10\,000$ K and $\log g$ shows a systematic across the whole $\log g$ range in the TAG problem at G=18.5. Comparing with the ILIUM results in Table 2, no very clear pattern emerges, with ILIUM superior to SVM in some problems and vice versa in others. When comparing performance star-by-star on a given problem, we do not see that one algorithm performs systematically better than the other in some parts of the AP space. Rather, one algorithm is just overall better. The fact that ILIUM is superior on the TG problems whereas SVM is superior on the TG-allmet problems suggests that SVM copes better with problems where there is more unmodelled variance, although the better performance of ILIUM on the TM-dwarfs problem and TAG at G=15 does not support this suggestion. Possibly it is when the variance is higher overall (lower SNR and/or unmodelled variance) that SVM gives smaller errors, which would be consistent with that method’s approach to dealing with noise. A more detailed comparison is only warranted after further work has been put into optimizing the ILIUM algorithm, such as the convergence for noisy data or the update clipping procedures.

SVM has the advantage of being much faster to apply once it has been tuned. As the speed of ILIUM is probably dominated by the nearest neighbour initialization, this could be replaced with an SVM. On the other hand, ILIUM provides AP uncertainty estimates, goodness-of-fit estimates (for outlier/poor solution detection) and allows one to inspect the relevance of each input in determining the output for every object (equation 8).