

The application of optimal estimation retrieval to lidar observations



Adam Charles Povey
St. John's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity Term 2013

Abstract

The application of optimal estimation retrieval to lidar observations

Adam Charles Povey, St. John's College

Submitted for the degree of Doctor of Philosophy in Atmospheric, Oceanic, and Planetary Physics in Trinity Term 2013

Optimal estimation retrieval is a nonlinear regression scheme to determine the conditions statistically most-likely to produce a given measurement, weighted against any *a priori* knowledge. The technique is applied to three problems within the field of lidar data analysis.

A retrieval of the aerosol backscatter and either the extinction or lidar ratio from two-channel Raman lidar data is developed using the lidar equations as a forward model. It produces profiles consistent with existing techniques at a resolution of 10–1000 m and uncertainty of 5–20 %, dependent on the quality of data. It is effective even when applied to noisy, daytime data but performs poorly in the presence of cloud.

Two of the most significant sources of uncertainty in that retrieval are the non-linearity of the detectors and the instrument's calibration (known as the dead time and overlap function). Attempts to retrieve a nonlinear correction from a pair of lidar profiles, one attenuated by a neutral density filter, are not successful as uncertainties in the forward model eliminate any information content in the measurements. The technique of Whiteman et al. [1992] is found to be the most accurate.

More successful is a retrieval of the overlap function of a Raman channel using a forward model combining an idealised extinction profile and an adaptation of the equations presented in Halldórsson and Langerholc [1978]. After refinement, the retrieval is shown to be at least as accurate, and often superior to, existing methods of calibration from routine measurements, presenting uncertainties of 5–15 %.

These techniques are then applied to observations of ash over southern England from the Eyjafjallajökull eruption of April 2010. Lidar ratios of 50–60 sr were observed when the plume first appeared, which reduced to 20–30 sr after several days within the planetary boundary layer, indicating an alteration of the particles over time.

Acknowledgements

This work was supported by a NERC CASE studentship with Hovemere Ltd. and would not have been possible without the support of all the staff there, especially Andrew Page, Andrey Delev, and David Rees. Judith Agnew's advice and repositories of raw data have been indispensable throughout this project. The continuous assistance, guidance, tolerance, and spare parts of Mark Munro, Russell Jones, and Duncan Constable have been a vital rectification of my own efforts. Without Warwick Clifton's blessing, RACHEL would never have made it on the roof and without the enthusiasm of David Howell, I'd never have gotten through the paperwork. I am indebted to Sarah Calcutt for juggling the ever-expanding beast that is my accounts. Thanks to everyone at the Chilbolton Observatory for their enthusiasm and extensive supply of spanners. I must acknowledge the important impact that gravity had on the practical aspects of this research, my own humility, and the balance sheet of Hovemere. The impeccable grammar and punctuation to follow is in great part due to the inexhaustible patience of my mother, though I claim full credit for any mistakes that remain. For a number of engaging conversations and keeping me on my toes, I thank Robin Hogan, Anthony Illingworth, Colin O'Connor, and everyone else in the radar-lidar group at the University of Reading. Special thanks are offered to Katherine Dunn, Katrin Thier, Naomi Petela, Emma Lewis, Glyn Kennington, Francis Travers, Amy Mason, Rhys Jones, and everyone else that have helped me maintain some semblance of sanity by sitting through my rants, sampling my cooking, or (to their almost certain detriment) both. Thanks to Dan Peters for his ability to not look annoyed when I stroll into his office. Finally, though this work is a long way from where he envisioned it, my supervisor Don Grainger was a foghorn over the shoals, keeping me from crashing into the shore and pointing towards a career ahead — my sincerest thanks for everything.

Contents

Abstract	i
List of Figures	x
List of Tables	xi
List of Acronyms	xii
List of Mathematical Notation	xiii
1 Introduction	1
1.1 Aerosol	2
1.2 Scattering	5
1.2.1 Molecular scattering	5
1.2.2 Particle scattering	6
1.2.3 Raman scattering	7
1.3 Light detection and ranging	8
1.3.1 Applications	8
1.3.2 The lidar equation	10
1.3.3 Photomultiplier tubes	14
1.4 Available instruments	16
1.4.1 RACHEL	16
1.4.2 CFARR	17
1.5 Nonlinear regression	18
1.5.1 Optimal estimation retrieval	18
1.5.2 Solution covariance	21
1.5.3 Convergence	22

2	Retrieval of aerosol extinction and backscatter	23
2.1	Existing techniques	23
2.1.1	Single channel lidar	23
2.1.2	Raman lidar	26
2.1.3	Regression methods	27
2.2	Forward model	29
2.3	Expected properties and distribution	33
2.3.1	Aerosol modelling	33
2.3.2	Simulated data	37
2.4	Configuration assessment	39
2.5	Sensitivity	45
2.6	Errors	50
2.6.1	Retrieval error	52
2.6.2	Parameter error	52
2.6.3	Further errors	56
2.7	Aside	57
2.8	Summary	59
3	Assessment of model parameters	61
3.1	Environmental terms	61
3.1.1	Pressure and temperature	61
3.1.2	Ångstrom coefficient	63
3.2	Instrumental effects	64
3.2.1	Vertical offset	64
3.2.2	Laser energy	65
3.2.3	Background counts	65
3.2.4	Detector nonlinearity	66
3.2.4.1	Forward model	68
3.2.4.2	Expected properties and distribution	69
3.2.4.3	Retrieval assessment	71
3.2.4.4	<i>Ad hoc</i> techniques	78
3.3	Summary	81

4	Estimation of the calibration function	83
4.1	Existing techniques	83
4.2	Forward model	85
4.2.1	Extinction	86
4.2.2	Overlap function	87
4.2.3	Summary	92
4.3	Expected properties and distribution	94
4.3.1	<i>A priori</i> assumptions	94
4.3.2	Distribution of state variables	95
4.3.3	Simulated data	99
4.4	Configuration assessment	100
4.5	Sensitivity	104
4.6	Errors	114
4.6.1	Retrieval error	114
4.6.2	Parameter error	119
4.7	Summary	120
5	Application to real data	123
5.1	Chilbolton ultraviolet lidar	123
5.1.1	Vertical offset	124
5.1.2	Measurement uncertainty	125
5.1.3	Calibration function	126
5.1.4	Extinction and backscatter	138
5.1.5	Eyjafjallajökull ash	151
5.2	RACHEL	155
6	Conclusions	161
A	Additional derivations	169
A.1	Lidar equation Jacobean matrix	169
A.1.1	Extinction configuration	169
A.1.2	Lidar ratio configuration	170
A.2	Lidar equation parameter errors	171
A.2.1	Dead time, τ_d	171
A.2.2	Background counts, E_B	171
A.2.3	Calibration function, C	171
A.2.4	Vertical offset, R_0	171

A.2.5	Ångstrom coefficient, κ	172
A.2.6	Number density, N	172
A.2.7	Molecular optical thickness, χ^m	172
A.3	Calibration function Jacobean matrix	172
A.3.1	Axis separation, d	173
A.3.2	Beam radius, w	173
A.3.3	Image size, ρ	174
A.3.4	Field of view, a	174
A.3.5	Telescope focusing, γ and ν	175
A.3.6	Laser pulse duration, τ_L	175
A.3.7	Optical efficiency, ζ	176
A.3.8	Optical thickness, χ	176
A.3.8.1	Exponential decay configuration	176
A.3.8.2	Error function configuration	176
A.4	Calibration function parameter errors	177
A.4.1	Dead time, τ_d	177
A.4.2	Range, R	177
A.5	Dead time Jacobean matrix	178
A.5.1	Nonparalyzable correction	178
A.5.2	Donovan correction	178
A.6	Simplification of ξ	178
A.7	Evaluation of S_H	181
A.8	IDL computer codes	184
A.8.1	Scattering forward model	185
A.8.2	Calibration function forward model	186
A.8.3	Effective area model	188
A.8.4	Optimal estimation retrieval	191
B	Additional tables and figures	195
B.1	Chapter 1	195
B.2	Chapter 2	196
B.3	Chapter 3	201
B.4	Chapter 4	202
B.5	Chapter 5	205

List of Figures

1.1	A sketch of the distribution of particle surface area, summarising sources and sinks	3
1.2	A sketch of the planetary boundary layer	4
1.3	A sketch of the spectrum of Raman scattering for molecular nitrogen	8
1.4	A schematic of lidar backscattering	11
1.5	The mechanism behind the lidar equation	14
1.6	Sketches of the operation of a PMT	15
1.7	Schematic of the optimal estimation retrieval algorithm	21
2.1	An application of the Fernald-Klett method	25
2.2	An application of the Ansmann method	26
2.3	The distribution of extinction for various aerosol models	34
2.4	<i>A priori</i> aerosol distributions	35
2.5	The observed vertical correlation of backscatter	37
2.6	A first application of the scattering retrieval	40
2.7	Attempts to improve the retrieval of cloudy data	41
2.8	Retrieval assuming a correlation between extinction and backscatter	42
2.9	Retrieval of the lidar ratio instead of extinction	43
2.10	Deviation of the retrieved from simulated profiles	44
2.11	Performance of the retrieval with noise-free data	46
2.12	The effect of increased SNR on the retrieval	47
2.13	Sensitivity to fine-scale fluctuations	49
2.14	A closer examination of fig. 2.13	50
2.15	Averaging kernels for fig. 2.13	51
2.16	Covariance matrix for retrieval of $\ln \beta$ and B from case 1	53
2.17	Retrieval in the lidar ratio configuration showing errors	54
2.18	As fig. 2.17, but for the extinction configuration	54
2.19	The contribution of model parameters to the variance	55
2.20	A demonstration of one-channel retrievals	58

3.1	An illustration of the variability in radiosonde profiles	63
3.2	Observed distribution of the 340–440 nm Ångstrom coefficient	64
3.3	Performance of the laser energy approximation	66
3.4	An example of laser energy variations in RACHEL	68
3.5	Covariance of analogue CUV observations	70
3.6	Performance of the dead time retrieval	72
3.7	Comparison of the dead time retrieved by three methods	73
3.8	As fig. 3.6, but comparing all three methods	74
3.9	As fig. 3.8, but only fitting points less than 15 MHz	75
3.10	Application of the Whiteman dead time estimation to RACHEL . . .	76
3.11	Pulse linearity of the Hamamatsu R7400U PMT	77
3.12	Nonlinearity observed in the CUV	79
3.13	The cost of fitting (3.2.9) to observations	79
3.14	Variations in $c_{0,1}$ from fitting (3.2.9)	80
4.1	An application of the Wandinger calibration method	85
4.2	A schematic of the optical axes of a lidar system	88
4.3	A geometric model of a Cassegrain telescope	89
4.4	Illustration that $\mathcal{H} \mu d\mu d\Theta = d\mathcal{A}(\mu, w; d)$	90
4.5	Illustration of the nine cases of S_H	91
4.6	Distribution of χ_∞ observed by the CFARR AERONET station	96
4.7	State vector PDFs for the erf configuration	97
4.8	As fig. 4.7, but for the exp configuration	97
4.9	The effect of Δ on the effective area	98
4.10	Components of C retrieved from simulated data	101
4.11	Residuals of the retrievals shown in fig. 4.10	102
4.12	Retrieved calibration function as a function of the SNR	103
4.13	The performance of various calibration function configurations . . .	103
4.14	CFARR’s overlap functions for the CUV	104
4.15	As fig. 4.12, but for cases 8–13	105
4.16	Perturbed aerosol profiles used for data simulation	106
4.17	Deviation of retrievals in the presence of a perturbed aerosol profile .	108
4.18	Deviation of retrievals of a Gaussian beam	109
4.19	Deviation of retrievals that used a randomised first guess	110
4.20	Deviation of retrievals as a function of Γ_0	111
4.21	Deviation of retrievals as a function of minimum retrieved height . .	111

4.22	The impact of annealing the retrieval	113
4.23	Confirmation of the error derived for the calibration function	115
4.24	Comparison of the retrieval and the Wandinger technique	116
4.25	Retrieved averaging kernel and covariance matrix	117
4.26	Contributions of model parameters to the variance	118
4.27	The impact of laser jitter on the calibration	119
5.1	An illustration of the vertical offset in the CUV	124
5.2	Variance of solar background as a function of its mean	126
5.3	Comparison of calibration functions retrieved for 21 Sep 2005	127
5.4	As fig. 5.3, but compared to the Wandinger method	129
5.5	As figures 5.3 and 5.4, but for 18 Jul 2007	130
5.6	Comparison of retrieved and Wandinger estimates for 2005–7	131
5.7	As fig. 5.6, but for 2 Mar 2010	132
5.8	As fig. 5.6, but for 16 Mar 2010	133
5.9	Calibration function retrieval errors	135
5.10	Comparison of overlap functions derived by three methods	136
5.11	Aerosol optical thickness derived from analogue profiles	137
5.12	Estimates of total backscatter and extinction for the test profiles	139
5.13	As fig. 5.12, but considering parameter errors during the retrieval	140
5.14	Averaging kernels for the retrievals of fig. 5.12	141
5.15	Errors for the retrieval of test case 10	142
5.16	Impact of C on the retrieved backscatter and extinction	143
5.17	Retrieved aerosol backscatter from 2 Mar 2010 showing errors	145
5.18	As fig. 5.17, but for the lidar ratio	146
5.19	A closer examination of figures 5.17 and 5.18	148
5.20	Supplementary observations from 2 Mar 2010	149
5.21	Retrieved χ_∞ compared to that observed by AERONET	149
5.22	Retrieval cost and laser energy for retrievals from 2 Mar 2010	151
5.23	Observations of the Eyjafjallajökull ash plume at CFARR	152
5.24	Distribution of extinction and backscatter for 19 Apr 2010	153
5.25	Retrieved overlap function for RACHEL	156
5.26	RACHEL observations of the Eyjafjallajökull ash plume	157
5.27	As fig. 5.26, but over Apr 17–19 2010	158
5.28	As fig. 5.24, but for RACHEL observations on 16 Apr 2010	159
A.1	The area of overlap of two circles	180

B.1	The theoretical Raman vibrational-rotational spectrum	195
B.2	Impact of vertical correlation scale height on the aerosol retrieval . .	197
B.3	Impact of vertical resolution on the aerosol retrieval	198
B.4	Covariance matrix for the lidar ratio configuration	199
B.5	Impact of minimum retrieved height on the aerosol retrieval	200
B.6	UK Met Office radiosonde launch locations in the UK	201
B.7	As fig. 4.7, but considering $\ln \delta, \zeta, H_{FT}$	203
B.8	As fig. 4.8, but considering $\ln \delta, \zeta, H_{FT}$	203
B.9	Deviation of retrievals as a function of maximum retrieved height . .	204
B.10	Vertical positioning of a cloud by various instruments	205
B.11	The failure of the measured laser energy on 2 Mar 2010	205
B.12	The data used to produce test profiles 1–5 for the C retrieval	208
B.13	As fig. B.12, but test profiles 6–12	209
B.14	As fig. 5.16, but showing errors	210
B.15	As fig. 5.16, but for the linear retrieval configuration	211

List of Tables

1.1	Vital statistics of the RACHEL system	17
1.2	Vital statistics of the CUV	19
2.1	<i>A priori</i> aerosol distributions for α/β retrieval	36
2.2	Details of simulated profiles	38
4.1	<i>A priori</i> state vector for the C retrieval	99
B.1	Additional details of simulated profiles for α/β retrieval	196
B.2	Dead time corrections of simulated profiles for α/β retrieval	196
B.3	Atmospheric parameters of simulated profiles for all retrievals	196
B.4	State vectors used to simulate data for C retrieval	202
B.5	Retrieved state vectors for photon counting data	206
B.6	Retrieved state vectors for analogue data	207

List of Acronyms

ABC	Attenuated backscatter coefficient
ADC	Analogue-to-digital converter
AERONET	Aerosol robotic network
BADC	British Atmospheric Data Centre
CALIPSO	Cloud-aerosol lidar and infrared pathfinder satellite observation
CFARR	Natural Environment Research Council Chilbolton Facility for Atmospheric and Radio Research
CUV	Chilbolton ultraviolet Raman lidar
DC	Direct current
DIAL	Differential absorption lidar
FOV	Field of view
IDL	Interactive Data Language
IPCC	Intergovernmental Panel on Climate Change
LIDAR	Light detection and ranging
NASA	National Aeronautics and Space Administration
NERC	National Environmental Research Council
NOAA	National Oceanic and Atmospheric Administration
ND	Neutral density
OPAC	Optical properties of aerosol and cloud
PBL	Planetary boundary layer
PC	Photon counting
PDF	Probability density function
PMT	Photomultiplier tube
RACHEL	Robust and compact hybrid environmental lidar
RADAR	Radio detection and ranging
RMSD	Root-mean square deviation
SNR	Signal-to-noise ratio
STFC	Science and Technology Facilities Council

List of Mathematical Notation

Symbol	Description	Units	Definition
$\alpha(\lambda, R)$	Aerosol extinction coefficient	km^{-1}	(1.2.6)
$\bar{\alpha}(\lambda, R)$	Volume extinction coefficient	km^{-1}	(1.2.9)
$\alpha^{\text{m}}(\lambda, R)$	Molecular extinction coefficient	km^{-1}	(1.2.5)
$\alpha_{BL,FT}$	Aerosol extinction within the planetary boundary layer or free troposphere, respectively	km^{-1}	(2.3.2)
α_{\pm}	$\frac{1}{2}(\alpha_{BL} \pm \alpha_{FT})$	km^{-1}	(4.2.3)
$\beta(\lambda^{\text{el}}, \lambda, R)$	Aerosol backscattering coefficient	$\text{Mm}^{-1} \text{sr}^{-1}$	(1.2.6)
$\bar{\beta}(\lambda^{\text{el}}, \lambda, R)$	Volume backscattering coefficient	$\text{Mm}^{-1} \text{sr}^{-1}$	(1.2.9)
$\beta^{\text{m}}(\lambda_L, \lambda, R)$	Molecular backscattering coefficient	$\text{Mm}^{-1} \text{sr}^{-1}$	(1.2.5)
$\beta^*(\lambda, R)$	Attenuated backscatter coefficient, $\bar{\beta}\mathcal{T}$	$\text{Mm}^{-1} \text{sr}^{-1}$	(1.3.20)
γ	$1 + \Delta/d$	—	(A.6.2)
δ	Distance between the telescope axis and laser pupil	m	(4.2.14)
δ_{ij}	The Kronecker delta	—	
ε	Calibration function parameters	various	
ϵ	Electric permittivity	F m	
ϵ	Measurement error	counts	
ζ^{ra}	$\langle E_L \rangle \eta^{\text{ra}} \frac{\lambda^{\text{el}}}{hc} \left \frac{d\sigma}{d\Omega} \right _{\lambda^{\text{el}}, \lambda^{\text{ra}}}$	$\text{counts m}^2 \text{sr}^{-1}$	(4.2.21)
$\eta(\lambda)$	Fraction of energy incident on the instrument at λ that is detected	—	
θ	Angular coordinate	rad	
κ	Ångstrom coefficient	—	(3.1.2)
λ	Wavelength, with superscripts el and ra indicating 354.7 nm and 386.7 nm, respectively. Functions of wavelength may use the shorthand $f(\lambda^{\text{el,ra}}) \equiv f^{\text{el,ra}}$	nm	
μ	Radial distance from the telescope's optical axis	m	Fig. 4.2

Symbol	Description	Units	Definition
$\nu(R)$	$ \gamma - \Delta R f^{-2} $	—	(A.6.7)
$\xi(R, \mathbf{r})$	Probability that a photon from a point \mathbf{r} on the surface a range R from the instrument reaches the detector, based on geometrical considerations	—	(4.2.9)
π	Ratio of the circumference of a circle to its diameter	—	
ρ_n	Depolarization factor	—	(1.2.3)
$\rho_{T,o}(R)$	Lens size in the scatterer plane, $\nu r_{T_o}/\gamma$	m	(4.2.10)
σ	Width of a Gaussian or error function	m	
$\sigma_R(\lambda)$	Rayleigh scattering cross-section	m ²	(1.2.4)
τ_b	Bin width, the duration over which the detectors integrate a measurement	ns	(3.2.1)
τ_d	Detector dead time	ns	
τ_L	Duration of laser pulse	ns	(1.3.9)
$\varphi(\lambda, R)$	Photon counts output by detectors, with subscripts m and t denoting measured and true, respectively	counts	
ϕ_L	Divergence of the laser beam	mrad	(4.2.12)
$\phi_{\parallel, \perp}$	Angle between the axes parallel or perpendicular to the plane defined by the telescope axis and laser pupil	mrad	(4.2.14)
$\chi(\lambda, R)$	Aerosol optical thickness, with the subscript ∞ denoting $R = \infty$	—	(2.2.5)
$\chi^m(R)$	Molecular optical thickness, with the subscript ∞ denoting $R = \infty$	—	
$\psi(R, \mathbf{r})$	Laser beam shape; the fraction of laser irradiance at a point \mathbf{r} on the surface a range R from the instrument	—	(1.3.3)
ω	Angular frequency	Hz	
Γ	Lapse rate	K km	
Γ_i	Scale constant in Levenberg-Marquadt iteration	—	(1.5.10)
Δ	Displacement of detectors from the focal plane	m	Fig. 4.3
Θ	Angular position in plane a distance R from the instrument	rad	Fig. 4.2
$\Upsilon(r_1, r_2; \mu)$	$[(r_1 + r_2)^2 - \mu^2][\mu^2 - (r_1 - r_2)^2]$	m ⁴	(A.6.17)

Symbol	Description	Units	Definition
$a(R)$	Aperture size in the scatterer plane, $\nu R r_a / \gamma f$	m	(4.2.10)
\mathbf{b}	Parameters of a forward model	various	
c	Speed of light in vacuum	m s^{-1}	
$d(R)$	Perpendicular distance between optical axes of the laser and telescope	m	Fig. 4.2
$dA(R, \mathbf{r})$	Unit area on the surface a range R from the instrument	m^2	
$\frac{d\sigma}{d\Omega}$	Differential cross-section for scattering	$\text{m}^2 \text{sr}^{-1}$	
f	Focal length of the telescope	m	Fig. 4.3
g	Gravitational acceleration	m s^{-1}	
$\mathbf{f}(\mathbf{x}, \mathbf{b}, \mathbf{b}')$	The unknown, exact forward model	counts	
h	Planck's constant	J s	
i	Bin number	—	
j	For a bin i , the subscript such that $r_j \leq R_i < r_{j+1}$	—	(2.2.7)
k	Wave-vector	m^{-1}	
k_B	Boltzmann's constant	J K ⁻¹	
l	The number of bins in r	—	
m	The number of bins in R	—	
m_A	Mean molecular mass of dry air	g	
n	Refractive index	—	
$p(R)$	Air pressure	hPa	
$q(\lambda, R, \mathbf{r})$	Probability that radiation emanating from a point \mathbf{r} on the surface a range R from the instrument will be detected	—	(1.3.5)
r	Range from instrument on the retrieval's vertical grid. Functions of this may be distinguished from functions of R by the superscript r	m	
\mathbf{r}	Position vector on surface a distance R from the instrument	m	
r_a	Radius of telescope aperture	m	Fig. 4.3
r_L	Initial radius of laser beam	m	(4.2.12)
r_o	Radius of telescope's secondary mirror	m	Fig. 4.3
r_T	Radius of telescope's primary mirror	m	Fig. 4.3

Symbol	Description	Units	Definition
t	Time; generally, the delay between pulse emission and observation	s	
$w(R)$	Half-width of the laser beam	m	Fig. 4.2
\mathbf{x}	State vector	various	
$\hat{\mathbf{x}}$	<i>A posteriori</i> solution	various	(1.5.10)
\mathbf{x}_a	<i>A priori</i> state vector	various	
\mathbf{y}	Measurement vector	various	(1.5.1)
z	Vertical height	m	
z_σ	$(z_{BL} - z)/\sigma_{BL}$	m	(4.2.3)
z_{BL}	Planetary boundary layer height	m	
z_H	$(z_{BL} - z)/H_{FT}$	m	(4.2.3)
$A(R, \epsilon)$	Effective area, $A_0 O(R)$	m ² sr	(4.2.22)
A_0	Area of the telescope's primary mirror	m ²	
$A_L(R)$	Laser cross-sectional area	m ²	(1.3.8)
\mathbf{A}	Averaging kernel, GK	various	(1.5.19)
$\mathcal{A}(r_1, r_2; \mu)$	Area of overlap between two circles, radii r_1 and r_2 , with distance μ between their centres	m ²	??
$\mathbb{A}(r_1, r_2; \mu)$	Area of overlap of two partially overlapped circles	m ²	(A.6.16)
A, B, C, D	Cubic spline interpolation matrices	—	(2.2.8)
$B(R)$	The lidar ratio, α/β	sr	(2.1.1)
\mathbb{B}	A convenience	m ⁴	(A.7.12)
\mathcal{B}	$8\pi/3$, the lidar ratio of Rayleigh scattering	sr	(1.2.5)
$C(\lambda, R)$	Calibration function	m ³ sr m ⁵	(1.3.15)
\mathbb{C}	A convenience	m ²	(A.7.22)
\mathcal{C}	Cost function	m ³ sr	(1.5.5)
D	Particle diameter	μm	(2.1.11)
$E(\lambda, R)$	Energy incident on detectors at wavelength λ within a bin at range R	J counts	(1.3.13)
$E_B(\lambda)$	Total background counts	counts	
E_L	Laser energy	J counts	
\mathbf{E}	Electric field	V m ⁻¹	
E, F	Cubic spline interpolation derivative matrices	m, m ⁻¹	(2.2.8)
$\mathbf{F}(\mathbf{x}, \mathbf{b})$	Forward model	various	

Symbol	Description	Units	Definition
\mathcal{F}	Relative power difference between profiles used to determine the dead time	—	(3.2.3)
$\mathbf{G}(\mathbf{x})$	Gain matrix, $\hat{\mathbf{S}} \mathbf{K}^T \mathbf{S}_\epsilon^{-1}$	various	
$\mathbf{G}_y(\mathbf{x})$	Sensitivity of the retrieved state to the measurement, $\partial \hat{\mathbf{x}} / \partial \mathbf{y}$	various	
H	A scale height	m	
\mathbf{H}	Magnetic field	T	
$\mathcal{H}(R)$	The Heaviside step function, which equals unity if its argument is positive and zero otherwise.	—	
$I(R)$	Laser irradiance	W m^{-2}	(1.3.9)
\mathbf{I}	Identity matrix	—	
$\mathbf{K}(\mathbf{x})$	Forward model Jacobean matrix, $\nabla_{\mathbf{x}} \mathbf{F}(\mathbf{x}, \mathbf{b})$	various	
$\mathbf{K}_b(\mathbf{x})$	Parameter Jacobean matrix, $\nabla_{\mathbf{b}} \mathbf{F}(\mathbf{x}, \mathbf{b})$	various	
$L_\lambda(\lambda, R, \mathbf{r})$	Laser-induced spectral radiance at a point \mathbf{r} on the surface a range R from the instrument	$\text{W m}^{-2} \text{sr}^{-1}$	(1.3.3)
M	Number of laser shots a profile is averaged over	—	
$\mathbf{M}, \mathbf{N}_{-mn}(kr, \theta, \phi)$	Vector spherical wave-functions of parity – and order mn	V m^{-1}	(1.2.7)
$N(R)$	Number density	m^{-3}	
\mathcal{N}	$\int_0^{R_i} N(R') \text{d}R'$	m^{-2}	(2.2.13)
$O(\lambda, R, \epsilon)$	Overlap function, which is parameterised by ϵ	sr	(1.3.8)
$P(\lambda, R)$	Instantaneous power at wavelength λ at a time t	W	(1.3.6)
$P(\mathbf{x})$	Probability of \mathbf{x} being true	—	
$P_{1,2,3}$	Coefficients of the Donovan dead time correction	—	(3.2.6)
$Q_{\text{ext},b}$	Mie scattering efficiency for extinction or backscattering, respectively	—	
R	Range from instrument on it's own vertical grid. Functions of range may use the shorthand $f(R_i) \equiv f_i$	m	(3.2.1)
R_0	Vertical offset of range axis	m	(3.2.1)

Symbol	Description	Units	Definition
R_c	A reference range at which the aerosol scattering is known	m	
$S(R)$	Corrected, log-elastic signal	—	(2.1.4)
$\mathcal{T}(\lambda, R)$	Transmission of the air column between the instrument and a range R at wavelength λ	—	
$\hat{\mathbf{S}}(\mathbf{x})$	Solution covariance matrix	various	(1.5.14)
\mathbf{S}_ϵ	Measurement covariance matrix	various	
\mathbf{S}_a	<i>A priori</i> covariance matrix	various	
\mathbf{S}_b	Parameter covariance matrix	various	
$S_H(\rho, R)$	Halldórsson S function, with superscripts indicating various cases of its solution	m ⁴	(4.2.23)
\mathbf{S}_y	Error covariance matrix	various	
$T(R)$	Air temperature	K	
$\langle \rangle$	Expectation value operator		

Chapter 1

Introduction

Aerosol is a suspension of solid or liquid particles in a gas. In the Earth's atmosphere, aerosols impact the radiation budget both directly, by reflecting solar radiation back into space [Haywood and Shine, 1995], and indirectly, by altering the properties and distributions of clouds [Lohmann and Feichter, 2005] or reacting with other species [Colbeck, 1998]. They present, at the time of writing, the single greatest uncertainty in estimates of net radiative forcing and therefore the ability to predict the impacts of climate change [IPCC [2007]].

Light detection and ranging, known as lidar, is an active remote sensing technique for observing the distribution of molecules and particles in the atmosphere as a function of height by means of the light they backscatter from a beam. The name intentionally emulates radar as both techniques use the time-of-flight of a pulsed source to deduce the distance to the scatterer, such that they return a profile of scattering against range from the instrument. Whilst radar uses radio waves to observe cloud droplets or large-scale variations in temperature, lidar employs wavelengths in the near-infrared to ultraviolet to sample the scattering of particles at unrivalled temporal and vertical resolutions [Fugii and Fukuchi, 2005].

Despite its strengths, lidar is not as widely applied as other measurements in studies of aerosol's direct and indirect effects. This is partially because a lidar presents a point measurement whilst a satellite provides global coverage. However, with the launch of a space-based lidar [Vaughan et al., 2004] and the development of lidar networks across Northern America [Welton et al., 2001], Europe [Pappalardo et al., 2005], and Asia [Sugimoto et al., 2008], there is an increasing volume of under-utilised data.

It is then necessary to consider the quality of data produced by lidars. In some quarters, there is a perception that lidar is a qualitative instrument — producing visually-appealing snapshots of aerosol events and height information, but not

rigorously-derived measurements of aerosol scattering or microphysical properties. The fact that the predominant techniques of lidar analysis were developed in the 1980s and in some institutions have not greatly evolved since then [Böckmann et al., 2004] in part justifies that belief. This thesis wishes to join the numerous current efforts to advance the field of lidar data analysis¹ to produce products more easily integrated with existing satellite retrievals, with a more rigorous assessment of their uncertainty, and using widely recognised retrieval algorithms that are less dependent on *ad hoc* corrections and assumptions.

The remainder of this chapter will introduce some fundamental concepts used throughout this thesis, including the technique of optimal estimation retrieval. Chapter 2 applies that technique to the estimation of aerosol extinction and backscatter from Raman lidar profiles and evaluates its performance with simulated data. Several practical concerns raised by that application will be addressed in Chapter 3, with Chapter 4 giving particular consideration to the development of a retrieval for the lidar calibration function — a measure of the sensitivity of the instrument as a function of range. All of these techniques will then be applied in Chapter 5 to actual observations from two Raman lidar systems, evaluating their performance against existing techniques and independent observations, and briefly considering some observations of the ash cloud from the Eyjafjallajökull eruption of April 2010. Chapter 6 discusses the conclusions of this research and outlines future work. Appendix A contains various derivations deemed unnecessary for an appreciation of the methodology, but included for reference, and Appendix B compiles extraneous figures and tables.

1.1 Aerosol

In the atmosphere, there are numerous types of aerosol from a variety of natural and anthropogenic sources including mineral dust blown from arid surfaces, sea spray, combustion products from fuel or biomass burning, and volcanic emissions such as ash or sulphur compounds [Seinfeld and Pandis, 1998]. These exhibit a wide range of optical properties due to variations in their shape, size, composition, and abundance [Bohren and Huffman, 2004]. Particles can also become incorporated into a water droplet or coated in other materials, further altering their optical properties. This plethora of different behaviours makes the remote identification

¹A few striking examples being Turner et al. [2002], Morille et al. [2007], Haeffelin et al. [2012], Mattis et al. [2012], and Pounder et al. [2012].

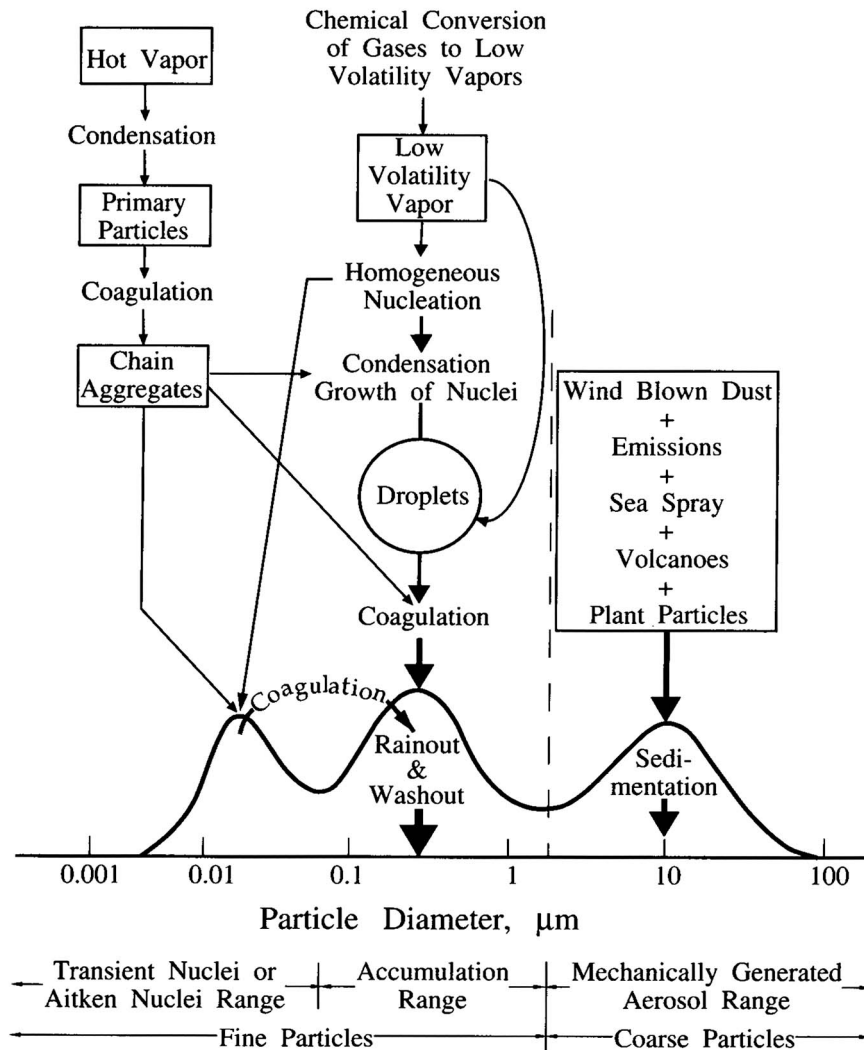


Figure 1.1: A sketch of the distribution of particle surface area for atmospheric aerosol including a summary of the sources and sinks for each mode. Originally figure 2.7 in Seinfeld and Pandis [1998], reproduced with permission. (©2006 by John Wiley & Sons, Inc. All rights reserved.)

and inventory of aerosol loading in the atmosphere a substantial observational challenge.

Aerosol diameters range from nanometres to tens of microns, with sub-micron particles often exceeding concentrations of 1000 cm^{-3} whilst larger particles are generally present in densities less than 1 cm^{-3} . The distribution of sizes, sketched in figure 1.1, can be accurately characterised by a log-normal distribution with three modes [Hinds, 1982]. The smallest is known as the nucleation mode as it mostly derives from gas-to-particle conversions in the atmosphere. Over time, these can coagulate into larger particles, producing the accumulation mode (which




Figure 1.2: A sketch of the planetary boundary layer over 24 hours, after Stull [1988]. (Image produced by Wikipedia user NikNaks, http://en.wikipedia.org/wiki/File:Atmospheric_boundary_layer.svg, accessed 5 Mar 2013, reproduced with permission.)

is frequently the most populous). Coarser aerosols ($> 2.5 \mu\text{m}$), corresponding to particulates, do not significantly interact with the other modes, having distinct sources and sinks.

Aerosols are removed from the atmosphere primarily by either dry deposition — the gravitational settling of particles to the surface — or wet deposition, where a particle becomes incorporated into a water droplet and leaves the atmosphere through precipitation. As these processes occur more rapidly than the time scale over which they would be mixed into the troposphere, the distribution of aerosols is highly inhomogeneous, often being concentrated in particular regions or around a singular event, such as desert dust over north Africa or a volcanic eruption. It is such variations that makes aerosol the greatest contribution to uncertainty in estimates of net radiative forcing.

The planetary boundary layer (PBL) is the lowest region of the atmosphere. Its dynamics are strongly influenced by surface friction and heating (in contrast to the free troposphere, where winds are mostly geostrophic) [Oke, 1987]. Within the PBL, there are three layers, sketched in figure 1.2:

- The surface layer, a thin region in direct contact with surface where drag is the primary driver of the flow;

- The mixed layer, a turbulent region extending over several metres to kilometres depending on the magnitude of surface heating that contains the greatest concentration of aerosol in the atmosphere as most emanate from the surface and are then dispersed through the layer by turbulence and convection, and;
- The entrainment layer, a thin temperature inversion capping the turbulent and convective updrafts of the mixed layer.

1.2 Scattering

1.2.1 Molecular scattering

Molecular, or Rayleigh, scattering corresponds to the classical solution for the elastic scattering of light by a single, small molecule. It is elastic in the sense that the wavelength of the light is not changed by the scattering and the molecule is small in comparison to that wavelength. By considering a plane, linearly-polarized wave of wavelength λ , it can be shown [Measures, 1992] that the cross-section for such scattering is,

$$\sigma_R(\lambda) = \frac{8\pi}{3} \frac{\pi^2 (n^2 - 1)^2}{N^2 \lambda^4}, \quad (1.2.1)$$

where n is the complex refractive index of the medium and N is the number density of scatterers. It can also be shown that the cross-section for backscattering (being scattering by an angle π) is,

$$\sigma_R(\lambda, \theta = \pi) = \frac{\pi^2 (n^2 - 1)^2}{N^2 \lambda^4}. \quad (1.2.2)$$

King [1923] showed that this result can be extended to account for the polarization of the wave,

$$\sigma_R(\lambda) = \frac{24\pi^3}{N^2 \lambda^4} \left(\frac{n^2 - 1}{n^2 + 2} \right)^2 \frac{6 + 3\rho_n}{6 - 7\rho_n}, \quad (1.2.3)$$

where ρ_n is the depolarization factor, which accounts for the anisotropy of the molecule. In practice, though, the atmosphere is composed of a mixture of molecules with various optical properties. Bucholtz [1995] reviewed measurements of a representative refractive index and depolarization for a ‘standard’ atmosphere² and derived the empirical fit,

$$\sigma_R(\lambda < 0.5 \mu\text{m}) = A\lambda^{-(B+C\lambda+D/\lambda)}, \quad (1.2.4)$$

²“Standard air is defined as dry air containing 0.03 % CO₂ by volume at normal pressure 760 mm Hg and having an air temperature of 15 °C.”

where λ is in microns with $A = 3.015, 77 \times 10^{-32} \text{ m}^2$, $B = 3.55212$, $C = 1.35579$, and $D = 0.11563$. This was shown to be consistent with observations to within 0.2 %.

The total extinction and backscattering suffered by a beam of light due to molecules in the atmosphere at some height z above the surface may be written as,

$$\alpha^m(\lambda, z) = \sigma_R(\lambda)N(z) \qquad \beta^m(\lambda, z) = \frac{\sigma_R(\lambda)N(z)}{\mathcal{B}}, \quad (1.2.5)$$

where $\mathcal{B} = \frac{8\pi}{3}$, the ratio of extinction to backscattering.

1.2.2 Particle scattering

Throughout this thesis, a two-component atmosphere will be considered, such that the total scattering (denoted with a bar) can be expressed as the sum of molecular and particle scattering,

$$\bar{\alpha} = \alpha + \alpha^m \qquad \bar{\beta} = \beta + \beta^m. \quad (1.2.6)$$

When the particle scattering is due to aerosols suspended in the atmosphere, it is common to estimate its magnitude with Mie theory [e.g. Wandering et al., 1995; Müller et al., 1999; Ferrare et al., 2001; Marchant et al., 2012]. Presented in Mie [1908], the theory is an exact formulation for the scattering of plane-polarized light by a spherical particle and so is most applicable to liquid or hydrated aerosols. Computer codes to calculate the scattering properties as a function of particle size and refractive index were developed some years ago and are easily obtained [e.g. Grainger et al., 2004].

For a more complete introduction to the method, consult Bohren and Huffman [2004], but to briefly summarise its main features — an x -polarized plane wave can be expressed in spherical coordinates by the vector spherical wave-functions $\mathbf{M}_{-mn}(kr, \theta, \phi)$ and $\mathbf{N}_{-mn}(kr, \theta, \phi)$, where ‘ $-$ ’ will denote even or odd functions (e or o , respectively) and m and n the order of that function with wave-vector $k = 2\pi/\lambda$. These are solutions to the wave equation $\nabla \mathbf{E} + k^2 \mathbf{E} = 0$ for the electric field \mathbf{E} . The incident electric field can then shown to be,

$$\mathbf{E}_i = E_0 \sum_{n=1}^{\infty} i^n \frac{2n+1}{n(n+1)} \left(\mathbf{M}_{o1n}^{(1)} - i \mathbf{N}_{e1n}^{(1)} \right). \quad (1.2.7)$$

The corresponding magnetic field can be obtained from $\nabla \times \mathbf{H} = -i\omega\epsilon\mathbf{E}$, where ω is the angular frequency of the wave and ϵ is the electric permittivity of the medium.

The scattered field \mathbf{E}_s and internal field \mathbf{E}_l can then be found by requiring that on the surface of the sphere, radius r_0 , the total field must disappear,

$$(\mathbf{E}_i + \mathbf{E}_s - \mathbf{E}_l) \times \hat{\mathbf{r}} = (\mathbf{H}_i + \mathbf{H}_s - \mathbf{H}_l) \times \hat{\mathbf{r}} = 0 \quad (r = r_0). \quad (1.2.8)$$

It can then be shown that the scattered field is a sum of Riccati-Bessel function and associated Legendre functions, from which a cross-section can be derived.

As aerosols are not uniformly sized, to derive the bulk scattering properties this cross-section must be integrated over particle size D ,

$$\alpha = \int \frac{\pi D^2}{4} Q_{\text{ext}} N(D) dD \quad \beta = \int \frac{\pi D^2}{4} Q_{\text{b}} N(D) dD, \quad (1.2.9)$$

where Q is the efficiency for a given process, being the cross-section for that process divided by the total cross-section for the particle to interact with the field in any way.

1.2.3 Raman scattering

Raman scattering is an inelastic scattering process whereby a photon excites a valence electron to a different rotational or vibrational energy level. As the levels are a function of the electronic configuration of the molecule, the Raman spectra is unique for every molecule. Figure 1.3 sketches the spectra resulting from scattering of green light by molecular nitrogen. The large, central peak is molecular scattering. The two next largest peaks correspond to a change in the electron's vibrational quantum number of ± 1 . About each of these peaks is the rotational spectrum, where the rotational quantum number also changes (more detail shown in figure B.1). Lines corresponding to a decrease in photon energy (increased λ) are known as Stokes lines, with anti-Stokes denoting the opposite.

The shape of the observed rotational spectra is a function of temperature, such that it can be used to sound atmospheric temperature [Zeyn et al., 1996], but mostly serves as a lower limit for the bandwidth of the molecular channel, which must capture enough of this spectrum to avoid temperature-dependent effects [Whiteman, 2003]. The backscattering coefficient will then be,

$$\beta(\lambda^{\text{ra}}, z) = N(z) \frac{d\sigma}{d\Omega}, \quad (1.2.10)$$

where $\frac{d\sigma}{d\Omega}$ is the differential cross-section for backscattering, a known constant³.

³For nitrogen, $3.5(1) \times 10^{-34} \text{ m}^2 \text{ sr}^{-1}$.




Figure 1.3: A sketch of the spectrum of Raman scattering, on a log scale, at 532 nm by molecular nitrogen. Reproduced with kind permission of the Elementary Processes in Gas Discharges (EPG) group, Eindhoven University of Technology, The Netherlands.

1.3 Light detection and ranging

1.3.1 Applications

Early studies photographed the beam of a searchlight as a means to estimate the density profile of the atmosphere from the Rayleigh scattering [Hulburt, 1937]. The invention of photomultiplier tubes gave improved vertical and temporal resolution [Friedland et al., 1956], but lidar techniques only began to be widely applied with the development of the Q-switched laser in early 1960s [Collis, 1966]. Its high energy and low beam divergence permit observations over a range of up to 100 km at a vertical resolution of around 10 m [Fugii and Fukuchi, 2005, chapter 1].

The intrinsic components of a lidar are fairly straightforward:

- A laser emitting into the atmosphere, possibly with some means to scan the beam. The preferred wavelength depends on the application whilst the pulse energy and repetition rate should be as high as safety requirements allow.
- A telescope to observe the backscattered light. This can be coincident with the laser beam or some distance away, depending on practical constraints.

- Some means to isolate the laser-induced signal from the background illumination, such as interference filters or an etalon.
- A photomultiplier tube to detect the signal, which is typically less than one photon per range bin, and the electronics required to gate and log those measurements.

As outlined to varying levels of detail in Measures [1992], Argall and Sica [2003], Fugii and Fukuchi [2005], and Weitkamp [2005], a number of variations of the basic technique have been developed for applications across the atmospheric sciences :

- *Ceilometers* are the simplest and most wide-spread lidar instrument, monitoring cloud base and PBL height [e.g. Wiegner et al., 2006; Haeffelin et al., 2012]. They use a single lens system to both transmit the laser beam and observe the backscattered light, which eliminates many calibration issues suffered by more complex systems. Low-powered lasers permit continuous, unattended operation and infrared wavelengths are exploited to minimise attenuation from Rayleigh scattering.
- *Rayleigh/Mie lidars* are the next simplest, profiling the elastic backscatter as a proxy of aerosol density. These use wavelengths of 500–1500 nm. However, their observations are limited to providing a relatively qualitative measure of aerosol abundance as there are several unknowns assessed by a single measurement.
- *Depolarising lidars* additionally monitor the polarization of the backscattered light as a means to estimate particle size or shape [e.g. Young, 1995; Marengo and Hogan, 2011]. CALIPSO is an example of such a system, currently flying within the NASA A-train [Vaughan et al., 2004].
- *Raman lidars* additionally monitor the $\Delta\nu = +1$ Raman scattering component⁴ of the spectrum. Ultraviolet wavelengths are used as the cross-section scales with λ^{-4} . As the most numerous constituent of the atmosphere, nitrogen is most commonly sounded, at a frequency shift of 2329.66 cm^{-1} [e.g. Melfi, 1972; Ansmann et al., 1992]. Water vapour, with a shift of 3651.7 cm^{-1} , can also be of interest [e.g. von Zahn et al., 2000; Guerrero-Rascado et al., 2008].

⁴This has a much larger cross-section than the anti-Stokes line as most molecules are in their ground state, though it is still $\sim 0.1 \%$ of the molecular signal.

- *Fluorescence lidar* is a less-common technique, sounding the density or temperature of a single species through the fluorescent emission of molecules excited by the laser beam [e.g. She et al., 1990].
- *Multiwavelength systems* use several lasers to sample the backscattering at a range of wavelengths, allowing estimation of microphysical parameters such as effective radius and aspect ratio [e.g. Althausen et al., 2000; Böckmann, 2001].
- *High spectral resolution lidar* discriminates between the Rayleigh and Mie backscattered light by measuring its spectrum and noting that Rayleigh scattering presents a broad, Doppler-broadened peak in contrast to a relatively narrow Mie peak (as aerosol particles have a larger mass) [e.g. Veselovskii et al., 2002; Hair et al., 2008].
- *Differential absorption lidar*, or DIAL, provides a means to sound more scarce species by using two closely spaced laser wavelengths — one on and one off of an absorption line of the species to be investigated [reviewed by Grant, 1991].
- *Doppler lidar* uses aerosols as a tracer to observe wind speeds through the PBL from the magnitude of the Doppler shift of the Rayleigh scattering peak [e.g. Hogan et al., 2009; Ghate et al., 2010].

1.3.2 The lidar equation

As illustrated in figure 1.4, a lidar directs a laser pulse into the atmosphere. As it travels, some energy is scattered out of the beam, the fraction of which that is within the telescope's field-of-view (FOV) can then be detected. The time delay between the emission of the pulse and its observation, t , gives the distance to the scattering site,

$$R = \frac{ct}{2}. \quad (1.3.1)$$

The lidar equation describes the energy observed from a range R as a function of the scattering properties of the atmosphere. The derivation that follows is adapted from Chapter 7 of Measures [1992] and is presented to introduce its fundamental assumptions and approximations.

Consider an infinitesimal volume of height dR and cross-sectional area $dA(R, \mathbf{r})$ centred about the point \mathbf{r} on the surface a range R from the instrument. The energy




Figure 1.4: A schematic of lidar backscattering where a volume of air scatters the incident laser energy, a fraction of which falls within the telescope's FOV and can be detected. Adapted from fig. 1.2 in Weitkamp [2005], used with permission. (©2005 by Springer. All rights reserved.)

eventually detected will depend on the laser-induced spectral radiance from that volume, L_λ , and the probability that radiation is observed by the detector, q . Integrating over the surface shows that the power incident due to a single laser shot from this range over the wavelength interval $(\lambda, \lambda + \Delta\lambda)$ is,

$$dP(\lambda, R) = dR \int L_\lambda(\lambda, R, \mathbf{r}) q(\lambda, R, \mathbf{r}) \Delta\lambda dA(R, \mathbf{r}). \quad (1.3.2)$$

The spectral radiance depends on two factors — the illumination of the beam and the fraction of it that is scattered towards the instrument. Hence,

$$L_\lambda(\lambda, R, \mathbf{r}) = \bar{\beta}(\lambda^{\text{el}}, \lambda, R, \mathbf{r}) I(R) \psi(R, \mathbf{r}), \quad (1.3.3)$$

where the laser has irradiance I and beam-shape ψ at wavelength λ^{el} ; and $\bar{\beta}$, known as the volume backscattering coefficient, is a sum across the scattering properties of all species in the atmosphere,

$$\bar{\beta}(\lambda^{\text{el}}, \lambda, R, \mathbf{r}) = \sum_X N_X(R, \mathbf{r}) \left| \frac{d\sigma}{d\Omega} \right|_{\pi, \lambda, X}, \quad (1.3.4)$$

where N_X is the number density of species X ; $\frac{d\sigma}{d\Omega}$ is the differential cross-section for scattering of a photon of wavelength λ^{el} towards the instrument with wavelength λ , which is assumed constant over the receiver's acceptance angle.

The probability of detection is a product of the fraction of light that isn't subsequently scattered out of the beam, known as the transmittance $\mathcal{T}(\lambda, R)$; the efficiency of the detectors and optical system, $\eta(\lambda)$; the acceptance solid angle of the

telescope (area A_0); and a factor to describe the probability that, considering the geometry of the optical system viewing a position (R, \mathbf{r}) , the photon could actually strike the detector, $\xi(R, \mathbf{r})$. Thus,

$$q(\lambda, R, \mathbf{r}) = \frac{A_0}{R^2} \mathcal{T}(\lambda, R) \eta(\lambda) \xi(\lambda, R, \mathbf{r}). \quad (1.3.5)$$

Integrating (1.3.2) over the range illuminated by the laser and using (1.3.3) with (1.3.5) gives,

$$\begin{aligned} P(\lambda, R) &= A_0 \int_0^R \frac{dR'}{R'^2} \int_{\Delta\lambda} d\lambda \eta(\lambda) \\ &\times \int \bar{\beta}(\lambda^{\text{el}}, \lambda, R', \mathbf{r}) I(R') \psi(R', \mathbf{r}) \mathcal{T}(\lambda, R') \xi(\lambda, R', \mathbf{r}) dA(R', \mathbf{r}). \end{aligned} \quad (1.3.6)$$

It is then assumed that the scattering medium is homogeneous over any given cross-section of the laser and that the spectral spread of both the scattered radiation and the laser is much smaller than the receiver's spectral window, such that variations in wavelength-dependent quantities can be ignored. Hence,

$$\begin{aligned} P(\lambda, R) &= A_0 \eta(\lambda) \int_0^R \frac{dR'}{R'^2} \bar{\beta}(\lambda^{\text{el}}, \lambda, R') I(R') \mathcal{T}(\lambda, R') \\ &\times \int \psi(R', \mathbf{r}) \xi(\lambda, R', \mathbf{r}) dA(R', \mathbf{r}). \end{aligned} \quad (1.3.7)$$

The second integral describes the overlap between the FOV of the telescope and the laser beam at a range R . This is a fairly elaborate function, which will be summarised by,

$$\int \psi(R, \mathbf{r}) \xi(\lambda, R, \mathbf{r}) dA(R, \mathbf{r}) = O(\lambda, R, \varepsilon) A_L(R), \quad (1.3.8)$$

where $O(\lambda, R, \varepsilon)$ is known as the overlap function, which will depend on some set of variables ε describing the instrument (discussed further in Chapter 4); and $A_L(R)$ is the cross-sectional area of the laser beam.

It is further assumed that the laser emits a short rectangular pulse of duration τ_L for which,

$$I(R) = \frac{E_L \mathcal{T}(\lambda^{\text{el}}, R)}{\tau_L A_L(R)}, \quad (1.3.9)$$

where E_L is the pulse energy. This changes the limits of the first integral in (1.3.7) to $[\frac{1}{2}c(t - \tau_L), \frac{1}{2}ct]$. If the pulse is short enough, the range-dependent parameters

may be treated as constants such that,

$$P(\lambda, R) = A_0 \eta(\lambda) O(\lambda, R, \epsilon) A_L(R) \bar{\beta}(\lambda^{\text{el}}, \lambda, R) I(R) \mathcal{T}(\lambda, R) \frac{2\tau_L}{ct(t - \tau_L)} \quad (1.3.10)$$

$$= A_0 \eta(\lambda) O(\lambda, R, \epsilon) \bar{\beta}(\lambda^{\text{el}}, \lambda, R) \mathcal{T}^* \frac{2E_L}{ct(t - \tau_L)}, \quad (1.3.11)$$

where $\mathcal{T}^* = \mathcal{T}(\lambda, R) \mathcal{T}(\lambda^{\text{el}}, R)$.

The detectors do not log the instantaneous power incident upon them, but the total energy over some interval $[t, t + \tau_b]$. Thus, a more relevant quantity is (assuming $\tau_b > \tau_L$ in order to observe the entire beam),

$$E(\lambda, R) = \int_{2R/c}^{2R/c + \tau_b} P(\lambda, R) dt \quad (1.3.12)$$

$$= A_0 \eta(\lambda) O(\lambda, R, \epsilon) \bar{\beta}(\lambda^{\text{el}}, \lambda, R) \mathcal{T}^* \frac{2E_L}{c\tau_L} \ln \left[\frac{R(R + \frac{c}{2}\tau_b - \frac{c}{2}\tau_L)}{(R + \frac{c}{2}\tau_b)(R - \frac{c}{2}\tau_L)} \right]. \quad (1.3.13)$$

In almost all cases this is simplified to,

$$E(\lambda, R) = \frac{E_L C(\lambda, R, \epsilon)}{R^2} \bar{\beta}(\lambda^{\text{el}}, \lambda, R) \mathcal{T}^* \quad (1.3.14)$$

$$C(\lambda, R, \epsilon) = A_0 \eta(\lambda) O(\lambda, R, \epsilon) \frac{2}{c\tau_L} R^2 \ln \left[\frac{R(R + \frac{c}{2}\tau_b - \frac{c}{2}\tau_L)}{(R + \frac{c}{2}\tau_b)(R - \frac{c}{2}\tau_L)} \right], \quad (1.3.15)$$

where C , which will be called the calibration function, gathers together the instrument-dependent terms. At almost all practically useful points $R \gg c\tau_L/2$ and $R \gg c\tau_b/2$, in which case the last two terms will approximately cancel giving,

$$C(\lambda, R, \epsilon) \approx A_0 \eta(\lambda) O(\lambda, R, \epsilon) \frac{c\tau_b}{2}. \quad (1.3.16)$$

Using the Beer-Lambert law, the transmission can be written in terms of a volume extinction coefficient, $\bar{\alpha}(\lambda, R)$, to obtain the basic lidar equation,


$$E(\lambda, R) = \frac{E_L C(\lambda, R, \epsilon)}{R^2} \bar{\beta}(\lambda^{\text{el}}, \lambda, R) \exp \left[- \int_0^R \bar{\alpha}(\lambda, R') + \bar{\alpha}(\lambda^{\text{el}}, R') dR' \right]. \quad (1.3.17)$$

In this thesis, two applications of this equation are considered:

- Elastic backscattering, where $\lambda = \lambda^{\text{el}}$,

$$E^{\text{el}}(R) = \frac{E_L C^{\text{el}}(R, \epsilon)}{R^2} \bar{\beta}^{\text{el}}(R) \exp \left[-2 \int_0^R \bar{\alpha}^{\text{el}}(R') dR' \right]. \quad (1.3.18)$$

Figure 1.5: The mechanism behind the lidar equation. **Left:** A lidar emits a pulse of energy E_L , which is attenuated by a factor \mathcal{T} en route to the scattering centre. **Right:** A fraction $\bar{\beta}$ of that is scattered towards the instrument. Again, a factor \mathcal{T} is attenuated en route, where the telescope observes a solid angle A_0/R^2 within that. Finally, a fraction C/A_0 of that light is focused onto the detectors and registered.



- Raman scattering, where λ is selected to sample the scattering from a single species, such that the sum of (1.3.4) can be ignored,

$$E^{\text{ra}}(R) = \frac{E_L C^{\text{ra}}(R, \epsilon)}{R^2} N(R) \exp \left[- \int_0^R \bar{\alpha}^{\text{el}}(R') + \bar{\alpha}^{\text{ra}}(R') dR' \right], \quad (1.3.19)$$

where the cross-section for Raman scattering has been subsumed into the calibration function.

These introduce a shorthand for functions of wavelength whereby,

$$f(\lambda^{\text{el,ra}}) \equiv f^{\text{el,ra}}.$$

A common lidar product is the attenuated backscatter coefficient (ABC),

$$\beta^*(R) = \frac{E^{\text{el}}(R) R^2}{E_L C^{\text{el}}(R, \epsilon)} \exp \left[2 \int_0^R \alpha^m(R') dR' \right] \quad (1.3.20)$$

$$\equiv \bar{\beta}(\lambda^{\text{el}}, \lambda, R) \exp \left[2 \int_0^R \alpha(R') dR' \right], \quad (1.3.21)$$

which is effectively a scattering measurement with all instrumental effects removed.

1.3.3 Photomultiplier tubes

Lidars use photomultiplier tubes (PMTs) as a detector⁵. These contain a photocathode — a material that is likely to emit an electron when struck by a photon. The

⁵Avalanche photodiodes are also widely used and their operation is broadly equivalent.

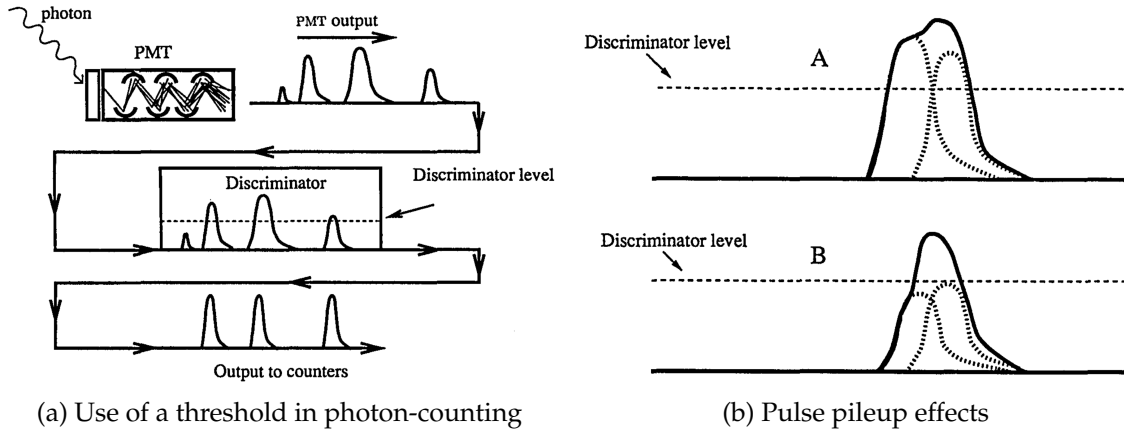


Figure 1.6: Figures 1 and 2 from Donovan et al. [1993], used with permission. (©1993 by OSA. All rights reserved.)

range of wavelengths that can be detected is bounded by the work function of the photocathode and the transmission of the material into which the cathode is set. That electron is accelerated towards a charged metal plate. Its impact releases further electrons, which are accelerated towards a second plate. This multiplication is repeated until the current is detectable. The shape of the current pulse represents the distribution of travel times through the tube (see sketches of fig. 1.6) and has a magnitude that is roughly Poisson distributed [Hakamata, 2006, chapter 6].

A PMT can be operated in one of two modes, depending on how the output is analysed. In the analogue mode, the output is logged by an analogue-to-digital converter (ADC), averaging the signal over some time period corresponding to a distance along the range axis. This behaves like an RC circuit to smooth the current pulse and is generally manufactured to be the dominant regulator of pulse shape. The output current is monitored and taken as proportional to the rate of photons striking the photocathode. When using a 12-bit ADC, a linear response can be obtained over up to three orders of magnitude, though this is reduced by DC offsets and is poorly constrained in its smallest bits. In addition to ambient light, thermal excitation in the photocathode and multiplier plates adds electrons to the system⁶, increasing the noise.

When the average number of counts in a bin is less than two (i.e. background illumination is minimal), a photon-counting mode can be applied. This notes that the majority of thermally excited cascades do not originate at the first multiplier plate nor possess as much kinetic energy as a photoelectron. Hence, these

⁶As do radioactive decays of radon within the detector's glass.

pulses have a lower magnitude. By only considering pulses that peak above some threshold (fig. 1.6a), the thermal background can be significantly reduced. By counting peaks rather than monitoring the current, the impact of the pulse height distribution can also be neglected.

However, the linearity of the detector is affected. As sketched in figure 1.6b, when two pulses arrive concurrently, they either (A) decrease or (B) increase the count rate by forming a single peak. This effect becomes increasingly important as the count rate increases and can be very complex, as extensively studied by the nuclear physics community [reviewed in Müller, 1973]. To avoid these difficulties, PMTs are generally designed such that the electronics are the greatest latency in the system and respond according to one of two extremes:

- Paralyzable detectors, where the detector is unable to observe another photon for a ‘dead’ time τ_d after any photon strike. The measured count rate φ_m can then be expressed in terms of the true count rate φ_t by,

$$\varphi_m = \varphi_t \exp(-\tau_d \varphi_t). \quad (1.3.22)$$

- Non-paralyzable detectors, where the detector is unable to observe another photon for a time τ_d after any *observed* photon strike, such that,

$$\varphi_m = \frac{\varphi_t}{1 + \tau_d \varphi_t}. \quad (1.3.23)$$

The latter is the most common in lidar studies [Whiteman, 2003, appendix B] and can frequently be used to correct signals up to 20 MHz.

Several lidar studies have also cited Donovan et al. [1993] as a nonlinear correction. This used an existing expression,

$$\varphi_m = \varphi_t \exp(-\tau_d \varphi_t) \left[P_1 + P_2 \tau_d \varphi_t + \frac{1}{2} P_3 (\tau_d \varphi_t)^2 + \dots \right], \quad (1.3.24)$$

where P_i represents the probability of i pulses combining into a peak that crosses the discriminator threshold that otherwise wouldn’t. That paper then described a means to estimate P_i as a function of the threshold voltage.

1.4 Available instruments

1.4.1 RACHEL

The Robust and Compact Hybrid Environmental Lidar (RACHEL), developed by Hovemere Ltd., is a cost-effective, four-channel, coaxial Raman lidar system designed for easy transport and unattended operation in the field. It includes the

Table 1.1: Vital statistics of the RACHEL system

<i>Laser</i>			
Model	Continuum Inlite III-20	Type	Nd-YAG
Wavelength, λ^{el}	354.7 nm	Beam expander	BXUV-10.0- 5X-355
Pulse rate	20 Hz	Mean pulse energy, $\langle E_L \rangle$	40 mJ
Transmitted beam diameter, $2r_L$	35 mm	Beam divergence, ϕ_L	0.3 mrad
Pulse length, τ_L	5–7 ns	Linewidth	1 cm ⁻¹
<i>Receiver</i>			
Model	Meade LXD-75	Efficiency	30 %
Primary mirror diameter, $2r_T$	203 mm	Secondary mirror diameter, $2r_o$	75 mm
Aperture diameter, $2r_a$	0.40 mm	Focal length, f	2 m
Elastic channel	355 nm	Raman channels	387, 407, 373 nm
<i>Detectors</i>			
Model	Hamamatsu H7155	Mode	Photon counting
Bin width, τ_b	60, 70 ns	Number of bins	4096 bins
Dead time (elastic), τ_d^{el}	40 ns	Dead time (Raman), τ_d^{ra}	70 ns
Efficiency	13 %	Dark count	50 Hz

capacity to scan the entire sky, intended as a means to monitor carbon dioxide emissions in factory plumes. Though that proved infeasible due to the minuscule carbon signal, it can readily applied to general pollution monitoring. Relevant details of the system are summarised in table 1.1.

1.4.2 CFARR

The majority of data to be evaluated in this thesis were collected at the Natural Environment Research Council Chilbolton Facility for Atmospheric and Radio Research (CFARR) [STFC [2011]] in Chilbolton, Hampshire (51.1445°, -1.4270°, 84 amsl). The site was founded in the 1960's as a test-bed for radio equipment — primarily

the 25 m 3 GHz Chilbolton Advanced Meteorological Radar. It currently hosts a wide variety of instrumentation for the observation of the atmosphere, including⁷:

- a 355 nm Leosphere EZ-lidar;
- a 905 nm Vaisala CT75K lidar ceilometer;
- a 1.5 μm Halo Photonics Doppler lidar;
- the Copernicus 35 GHz cloud radar;
- an AERONET station [Holben et al., 1998; Woodhouse and Agnew, 2011];
- daily radiosonde launches from Larkhill, approximately 30 km west of the site [UK Met Office [2011]];
- micro and millimetre-wave, visible, and infrared radiometers, and;
- meteorological instrumentation with a cloud camera.

Of greatest interest to this work, though, is the Chilbolton Ultraviolet Raman lidar (CUV) [Agnew, 2003; Agnew and Wrench, 2010]. It is optimised for Raman water vapour profiling through the daytime boundary layer and so implements both photon counting and analogue data collection, with a concentration on the latter due to its minimised background noise during the day. RACHEL was stationed at CFARR from January to May 2010 for an intercomparison with the CUV as the systems monitor identical Raman lines, though RACHEL samples a much lower power due to its smaller telescope and laser. Direct comparisons can also be made with the EZ lidar, which provides depolarization profiles instead of Raman observations. The other radar and lidar products can also provide some less direct comparisons to thoroughly investigate any discrepancies found and assess the performance of the RACHEL system.

1.5 Nonlinear regression

1.5.1 Optimal estimation retrieval

Optimal estimation is a form of nonlinear regression intend to solve the inverse problem,

$$\mathbf{y} = \mathbf{F}(\mathbf{x}, \mathbf{b}) + \boldsymbol{\epsilon}, \quad (1.5.1)$$

⁷Full specifications and data available at <http://badc.nerc.ac.uk/data/chilbolton/>. Current observations can be found at <http://www.stfc.ac.uk/chilbolton/26688.aspx>.

Table 1.2: Vital statistics of the CUV [Agnew, 2002]

<i>Laser</i>			
Model	Continuum PL9050	Type	Nd-YAG
Wavelength, λ^{el}	354.7 nm	Beam expander	$\times 10$
Pulse rate	50 Hz	Mean pulse en- ergy, $\langle E_L \rangle$	350 mJ
Transmitted beam dia- meter, $2r_L$	100 mm	Beam divergence, ϕ_L	0.1 mrad
Pulse length, τ_L	7 ns	Linewidth	1 cm^{-1}
<i>Receiver</i>			
Model	Newtonian	Focal length, f	2 m
Primary mirror dia- meter, $2r_T$	460 mm	Secondary mirror diameter, $2r_o$	100 mm
Aperture diameter, $2r_a$	8 mm		
<i>Detectors</i>			
Model	Hamamatsu R7400U-03	Mode	Both
Bin width, τ_b	50 ns	Number of bins	1800 bins
Anode sensitivity	50 A lm^{-1}	Dark count	80 Hz
Elastic channel	354.7 nm	Raman channels	386.7, 407.7, 353.0, 353.9 nm

for the the state vector, \mathbf{x} , as a function of the measurement vector, \mathbf{y} , through a forward model, $\mathbf{F}(\mathbf{x}, \mathbf{b})$, dependent on known parameters, \mathbf{b} , with a measurement error ϵ . For a lidar, the measurement vector would describe the number of photons observed in each height bin for each channel; the state vector describes the unknown scattering properties of the atmosphere; and the forward model is an equation describing how the instrument and light interact with the atmosphere and each other to produce a measurement [i.e. the lidar equation (1.3.17)].

The solution $\hat{\mathbf{x}}$ is the state with the maximal probability of representing the true state of the atmosphere given the measurements (known as the maximum *a posteriori* solution). It can be found using Bayesian statistics and assuming the probability density functions (PDFs) describing all quantities are Gaussian (or at least symmetric). In such circumstances,

$$-2 \ln P(\mathbf{y}|\mathbf{x}) = [\mathbf{y} - \mathbf{F}(\mathbf{x}, \mathbf{b})]^T \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{F}(\mathbf{x}, \mathbf{b})] + c_1, \quad (1.5.2)$$

where c_1 is a constant and $\mathbf{S}_\epsilon^{-1} = \langle [\mathbf{y} - \bar{\mathbf{y}}][\mathbf{y} - \bar{\mathbf{y}}]^T \rangle$ is the measurement error covariance. The prior knowledge of the system is summarised by an *a priori* state vector \mathbf{x}_a which, though not necessarily realistic, is also modelled with a symmetric probability density function such that,

$$-2 \ln P(\mathbf{x}) = [\mathbf{x} - \mathbf{x}_a]^T \mathbf{S}_a^{-1} [\mathbf{x} - \mathbf{x}_a] + c_2, \quad (1.5.3)$$

where \mathbf{S}_a^{-1} is the covariance matrix of the *a priori*.

Substituting these into Bayes' Theorem,

$$P(\mathbf{x}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{x})P(\mathbf{x})}{P(\mathbf{y})}, \quad (1.5.4)$$

gives the cost function,

$$\mathcal{C} = -2 \ln P(\mathbf{x}|\mathbf{y}) = [\mathbf{y} - \mathbf{F}(\mathbf{x}, \mathbf{b})]^T \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{F}(\mathbf{x}, \mathbf{b})] + [\mathbf{x} - \mathbf{x}_a]^T \mathbf{S}_a^{-1} [\mathbf{x} - \mathbf{x}_a], \quad (1.5.5)$$

where it has been assumed that $P(\mathbf{y})$ is a normalising constant and therefore neglected. Practically, the first term is similar to the root-mean square deviation (RMSD) of the modelled measurement from that observed, weighted by the error on that observation. The second term is then the deviation of the state from the *a priori*, weighted by its expected distribution.

The maximum probability state $\hat{\mathbf{x}}$ then solves,

$$\begin{aligned} 0 &= \nabla_{\mathbf{x}} \{-2 \ln[P(\hat{\mathbf{x}}|\mathbf{y})]\} \\ &= -\mathbf{K}^T \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{F}(\hat{\mathbf{x}}, \mathbf{b})] + \mathbf{S}_a^{-1} [\hat{\mathbf{x}} - \mathbf{x}_a]. \end{aligned} \quad (1.5.6)$$

where $\mathbf{K} = \nabla_{\mathbf{x}} \mathbf{F}(\hat{\mathbf{x}}, \mathbf{b})$.

The vector equation $\mathbf{g}(\mathbf{x}) = 0$ can be solved with the Newton iteration,

$$\mathbf{x}_{i+1} = \mathbf{x}_i - [\nabla_{\mathbf{x}} \mathbf{g}(\mathbf{x}_i)]^{-1} \mathbf{g}(\mathbf{x}_i), \quad (1.5.7)$$

which converges to $\hat{\mathbf{x}}$. To solve (1.5.6),

$$\mathbf{g} = \mathbf{S}_a^{-1} [\mathbf{x}_i - \mathbf{x}_a] - \mathbf{K}^T \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{F}(\mathbf{x}_i, \mathbf{b})] \quad (1.5.8)$$

$$\nabla_{\mathbf{x}} \mathbf{g} = \mathbf{S}_a^{-1} + \mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K} - [\nabla_{\mathbf{x}} \mathbf{K}^T] \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{F}(\mathbf{x}_i, \mathbf{b})]. \quad (1.5.9)$$

In the quasi-linear case, the last term of (1.5.9) is neglected to give the Gauss-Newton iteration. However, if the forward model is more nonlinear, this term can instead be approximated by $\Gamma_i \mathbf{S}_a^{-1}$, where Γ_i is chosen at each step to reduce the cost function (1.5.5). This is known as the Levenberg-Marquardt iteration,

$$\begin{aligned} \mathbf{x}_{i+1} &= \mathbf{x}_i + [(\mathbf{1} + \Gamma_i) \mathbf{S}_a^{-1} + \mathbf{K}_i^T \mathbf{S}_\epsilon^{-1} \mathbf{K}_i]^{-1} \\ &\quad \{\mathbf{K}_i^T \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{F}(\mathbf{x}_i, \mathbf{b})] - \mathbf{S}_a^{-1} (\mathbf{x}_i - \mathbf{x}_a)\}. \end{aligned} \quad (1.5.10)$$

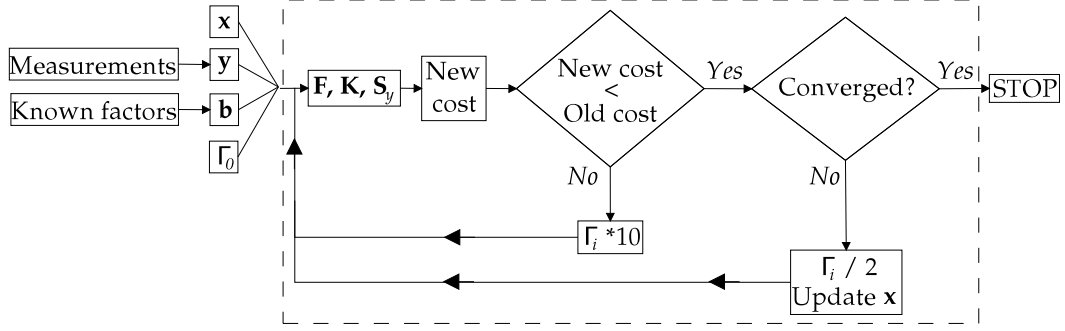


Figure 1.7: Schematic of optimal estimation retrieval algorithm.

1.5.2 Solution covariance

The covariance of the solution can be estimated by linearising the forward model about it, $F(\hat{\mathbf{x}}, \mathbf{b}) \approx \hat{\mathbf{K}}\hat{\mathbf{x}}$, such that the cost function can be written as,

$$-2 \ln P(\mathbf{x}|\mathbf{y}) = [\mathbf{y} - \mathbf{K}\mathbf{x}]^T \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{K}\mathbf{x}] + [\mathbf{x} - \mathbf{x}_a]^T \mathbf{S}_a^{-1} [\mathbf{x} - \mathbf{x}_a] \quad (1.5.11)$$

$$\equiv (\mathbf{x} - \hat{\mathbf{x}})^T \hat{\mathbf{S}}^{-1} (\mathbf{x} - \hat{\mathbf{x}}), \quad (1.5.12)$$

where the second line is simply an alternative factorisation of the quadratic.

Equating terms quadratic in \mathbf{x} ,

$$\mathbf{x}^T \hat{\mathbf{S}}^{-1} \mathbf{x} = \mathbf{x}^T \mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K} \mathbf{x} + \mathbf{x}^T \mathbf{S}_a^{-1} \mathbf{x} \quad (1.5.13)$$

$$\therefore \hat{\mathbf{S}} = (\mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K} + \mathbf{S}_a^{-1})^{-1}. \quad (1.5.14)$$

Then, considering terms linear in \mathbf{x}^T ,

$$-\mathbf{x}^T \hat{\mathbf{S}}^{-1} \hat{\mathbf{x}} = -\mathbf{x}^T \mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{y} - \mathbf{x}^T \mathbf{S}_a^{-1} \mathbf{x}_a \quad (1.5.15)$$

$$\hat{\mathbf{x}} = \hat{\mathbf{S}} [\mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{y} + (\hat{\mathbf{S}}^{-1} - \mathbf{K}^T \mathbf{S}_\epsilon^{-1} \mathbf{K}) \mathbf{x}_a] \quad (1.5.16)$$

$$= \mathbf{x}_a + \hat{\mathbf{S}} \mathbf{K}^T \mathbf{S}_\epsilon^{-1} (\mathbf{y} - \mathbf{K} \mathbf{x}_a). \quad (1.5.17)$$

Defining the gain matrix $\mathbf{G} = \hat{\mathbf{S}} \mathbf{K}^T \mathbf{S}_\epsilon^{-1}$ and using (1.5.1),

$$\hat{\mathbf{x}} = \mathbf{x}_a + \mathbf{G} (\mathbf{K} \mathbf{x} + \epsilon - \mathbf{K} \mathbf{x}_a) \quad (1.5.18)$$

$$= (\mathbf{I} - \mathbf{A}) \mathbf{x}_a + \mathbf{A} \mathbf{x} + \mathbf{G} \epsilon, \quad (1.5.19)$$

where $\mathbf{A} = \mathbf{G} \mathbf{K}$, the averaging kernel. It describes the extent to which the true and *a priori* states each contribute to the solution. An ideal retrieval would have a kernel equal to the identity, whilst practical kernels can be used to assess the resolution of the retrieval as (if \mathbf{x} describes a profile) the rows of \mathbf{A} will generally

be peaked functions showing how the information in one retrieved bin is derived from an average of those around it. Further, the trace of the averaging kernel is a measure of the number of degrees of freedom for signal in the retrieval [Rodgers, 2000, chapter 2.4.2].

1.5.3 Convergence

General practice, outlined in figure 1.7, is that if the cost function increases after an iteration, Γ_i is increased by a factor of ten. Otherwise, it is reduced by a factor of two.

The algorithm converges when either:

- The cost function decreases by less than some fraction of the number of measurements;
- The cost decreased and the change in the state is less than some fraction of the predicted error,

$$(\mathbf{x}_{i+1} - \mathbf{x}_i)_j < \epsilon \sqrt{\hat{\mathbf{S}}_{jj}} \quad \forall j. \quad (1.5.20)$$

- The step in state space,

$$\{\mathbf{K}_i^T \mathbf{S}_\epsilon^{-1} [\mathbf{y} - \mathbf{F}(\hat{\mathbf{x}}_i, \mathbf{b})] - \mathbf{S}_a^{-1} (\hat{\mathbf{x}}_i - \mathbf{x}_a)\} (\mathbf{x}_{i+1} - \mathbf{x}_i)^T, \quad (1.5.21)$$

is less than some fraction of the length of the state vector.

Otherwise, evaluation will cease after 30 iterations, which is considered a failure.

The goodness of the fit can be judged by considering the value of the cost function once converged. Solutions that are good models should have a cost approximately equal to the number of measurements. For this reason, when quoted, the cost will always be normalised by the length of \mathbf{y} , such that good retrievals should converge to a cost of approximately unity.

Chapter 2

Retrieval of aerosol extinction and backscatter

The volume extinction and backscattering coefficients are functions of more fundamental, microphysical properties of the scatterers, such as refractive index and effective radius. Deriving these properties directly is possible and of scientific interest, but the problem is poorly constrained. Its solution either requires a greater number of measurements, such as a multi-wavelength system [Müller et al., 1999; Marchant et al., 2012] or further assumptions about the scatterers, as in the observation of water clouds where the particle refractive index and size distribution can be reasonably approximated [Delanoe and Hogan, 2008]. These complexities are frequently disregarded in favour of attempting the better-constrained retrieval of extinction and backscatter.

2.1 Existing techniques

2.1.1 Single channel lidar

Single channel lidar clearly presents an underconstrained problem. However, it is by far the simplest system to operate and, as any lidar can be analysed as a single channel instrument, can provide a first guess for more involved analyses (e.g. Wandinger and Ansmann [2002]). Further, an increasing number of ceilometers are being adapted to log the scattering profile, widening the range of applications for single channel analyses [Mattis et al., 2012].

The most widely used algorithm is referred to as the Fernald-Klett method or onion peeling, developed across Fernald et al. [1972], Klett [1981], Fernald [1984],

and Klett [1985]. The (aerosol) lidar ratio B ,

$$\alpha(R) = B(R) \beta(R), \quad (2.1.1)$$

is assumed known such that the lidar equation (1.3.18) can be rearranged,

$$\frac{E^{\text{el}}(R) R^2}{C^{\text{el}}(R, \epsilon)} = E_L \bar{\beta}(R) \exp \left[-2 \int_0^R \mathcal{B} \beta^{\text{m}}(R') + B(R') \beta(R') \, dR' \right] \quad (2.1.2)$$

$$= E_L \bar{\beta}(R) \exp \left[-2 \int_0^R \mathcal{B} \beta^{\text{m}}(R') + B(R') [\bar{\beta}(R') - \beta^{\text{m}}(R')] \, dR' \right], \quad (2.1.3)$$

where $\mathcal{B} = \frac{8\pi}{3}$, the lidar ratio for Rayleigh scattering.

It is then possible to define a ‘corrected’ signal,

$$S(R) = \ln \left[\frac{E^{\text{el}}(R) R^2}{C^{\text{el}}(R, \epsilon)} \right] + 2 \int_0^R [\mathcal{B} - B(R')] \beta^{\text{m}}(R') \, dR' \quad (2.1.4)$$

$$\equiv \ln E_L + \ln \bar{\beta}(R) - 2 \int_0^R B(R') \bar{\beta}(R') \, dR'. \quad (2.1.5)$$

Differentiating,

$$\frac{dS}{dR} = \frac{1}{\bar{\beta}(R)} \frac{d\bar{\beta}}{dR} - 2B(R) \bar{\beta}(R), \quad (2.1.6)$$

and making the substitution $x = \bar{\beta}^{-1}$,

$$\frac{dx}{dR} + x \frac{dS}{dR} = -2B(R). \quad (2.1.7)$$

Solving this by means of an integrating factor, integrating from R to a range where the scattering properties are known¹ R_c ,

$$x_c e^{S(R_c)} - x e^{S(R)} = -2 \int_R^{R_c} B(R') e^{S(R')} \, dR' \quad (2.1.8)$$

$$\begin{aligned} \therefore \bar{\beta}(R) &= \frac{e^{S(R)}}{\exp^{S(R_c)} [\bar{\beta}(R_c)]^{-1} + 2 \int_R^{R_c} B(R') e^{S(R')} \, dR'} \\ &\equiv \frac{E^{\text{el}}(R) R^2 \exp \left[-2 \int_R^{R_c} [\mathcal{B} - B(R')] \beta^{\text{m}}(R') \, dR' \right]}{\frac{E^{\text{el}}(R_c) R_c^2}{\bar{\beta}(R_c)} + 2 \int_R^{R_c} \frac{B(R') E^{\text{el}}(R') R'^2}{\exp \left[2 \int_{R'}^{R_c} [\mathcal{B} - B(R'')] \beta^{\text{m}}(R'') \, dR'' \right]} \, dR'}. \end{aligned} \quad (2.1.9)$$

An iterative scheme is then developed by using the trapezium rule to estimate $\bar{\beta}$ one bin below the reference height,

$$\bar{\beta}_{c-1} = \frac{E_{c-1}^{\text{el}} R_{c-1}^2 e^D}{E_c^{\text{el}} R_c^2 \bar{\beta}_c^{-1} + (B_c E_c^{\text{el}} R_c^2 + B_{c-1} E_{c-1}^{\text{el}} R_{c-1}^2 e^D)^{\frac{1}{2}} c\tau_b} \quad (2.1.10)$$

$$D = \frac{1}{2} c\tau_b [(B_c - \mathcal{B}) \beta_c^{\text{m}} + (B_{c-1} - \mathcal{B}) \beta_{c-1}^{\text{m}}], \quad (2.1.11)$$

$$(2.1.12)$$

¹Integrating from the surface is unstable as errors increase exponentially [Bissonnette, 1986].

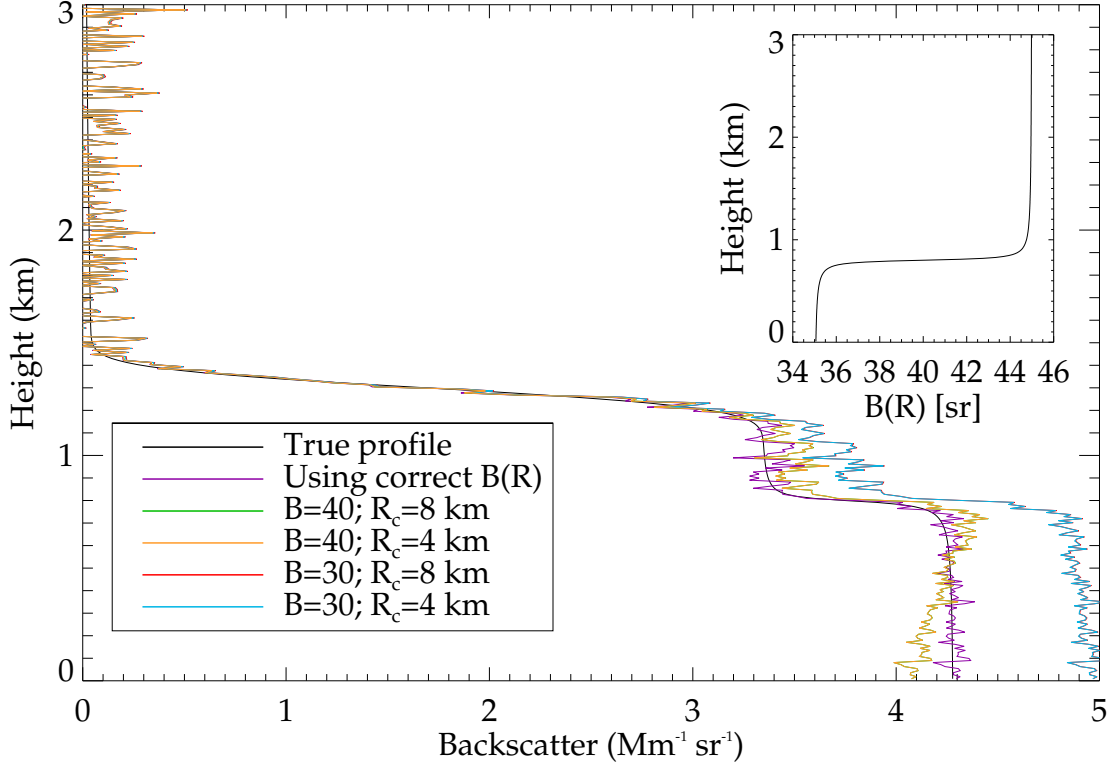


Figure 2.1: An application of the Fernald-Klett method to data simulated for 144,000 20 mJ shots, showing a dependence on the assumed lidar ratio B and reference height R_c . The simulated lidar ratio profile is inset.

where $\frac{1}{2}c\tau_b$ is the bin width. This introduces a notation that will be used throughout this thesis, where the subscript denotes range bin number. As the range axis is fundamentally a discrete series of bins, when processing data the bin number is a more practical quantity to work with. Hence, for all range dependent variables,

$$f_i \equiv f(R_i). \quad (2.1.13)$$

The complete profile is determined by evaluating (2.1.10) for each step down from the reference height in turn, effectively using each result as the reference for the next. These will clearly be affected by noise, but compared to a ‘proper’ integration of (2.1.9), the impact is reduced. However, it is dependent on the initial reference value E_c^{el} . Herein, that will be determined by evaluating $\bar{\beta}(R)$ for a range of points believed to clear of aerosol and minimising the absolute value of $\langle \bar{\beta} - \beta^m \rangle$ over that range as a function of E_c^{el} .

An implementation of this algorithm to simulated data is shown in figure 2.1, assuming lidar ratios of 30 or 40 sr. Different reference heights are also considered.




Figure 2.2: Application of the Ansmann method to data (inset) simulated for 144,000 20 mJ shots. Linear fits to 200 m of data are overplotted on (2.1.15) with the resulting aerosol extinction plotted (red) against the true profile (right).

It is clear that errors in the lidar ratio dominate whilst the choice of reference height is of minimal importance (provided it is actually clean). If measurements of χ_∞ are available, a lidar ratio can be selected to give an extinction profile consistent with that measurement. Recent implementations of this algorithm include pre-processing to identify aerosol and cloud layers within the profile, from which the lidar ratio profile is estimated either by *a priori* assumption of the type of aerosol present or consideration of variations in the integrated backscatter as a function of polarization [Vaughan et al., 2004; Böckmann et al., 2004; Cao et al., 2009]

2.1.2 Raman lidar

The Raman lidar equation (1.3.19) is a function of only extinction (if the number density profile is known). Its most common method of solution was presented in Ansmann et al. [1990] by assuming an Ångström coefficient of unity²,

$$\frac{\alpha(\lambda^{\text{ra}}, R)}{\alpha(\lambda^{\text{el}}, R)} = \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}}. \quad (2.1.14)$$

²That paper also considered a value of zero for ice particles, but as those will not be considered in this thesis, that distinction is neglected.

Rearranging (1.3.19) and differentiating with respect to R ,

$$\frac{d}{dR} \left[\ln \frac{N(R) C^{\text{ra}}(R)}{E^{\text{ra}}(R) R^2} \right] = \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \alpha(\lambda^{\text{el}}, R) + \alpha^{\text{m}}(\lambda^{\text{el}}, R) + \alpha^{\text{m}}(\lambda^{\text{ra}}, R) \quad (2.1.15)$$

$$\therefore \alpha(\lambda^{\text{el}}, R) = \frac{\frac{d}{dR} \left[\ln \frac{N(R) C^{\text{ra}}(R)}{E^{\text{ra}}(\lambda, R) R^2} \right] - \alpha^{\text{m}}(\lambda^{\text{el}}, R) - \alpha^{\text{m}}(\lambda^{\text{ra}}, R)}{1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}}}. \quad (2.1.16)$$

Evaluation of the derivative is the most difficult aspect of this technique, as even low noise levels can dominate the gradient. The simplest means to minimise this is smoothing the data in time and space to increase the signal-to-noise ratio (SNR) [Ansmann et al., 1992]. Alternatively, a linear fit to the processed data can be performed, from which the derivative can be easily extracted [Whiteman, 1999; Turner et al., 2002]. This is illustrated in figure 2.2, highlighting how the method can produce negative α , even over long averaging periods.

The backscatter can be derived by substituting the derived extinction into the elastic lidar equation and rearranging. Alternatively, the paper proposed making use of a clear reference height such that,

$$\beta_i^{\text{el}} = \beta_i^{\text{m}}(\lambda_0) \left\{ -1 + \frac{E_c^{\text{ra}} E_i^{\text{el}}}{E_i^{\text{ra}} E_c^{\text{el}}} \exp \left[\int_{R_i}^{R_c} \bar{\alpha}^{\text{ra}}(R) - \bar{\alpha}^{\text{el}}(R) dR \right] \right\}, \quad (2.1.17)$$

which will be referred to as the Ansmann ratio.

2.1.3 Regression methods

For several decades, the technique of optimal estimation has been successfully applied to the analysis of satellite [Marks and Rodgers, 1993; Palmer et al., 2000; Li et al., 2008; Watts et al., 2011] and ground-based radiometer observations [Guldner and Spankuch, 2001; Haefele et al., 2012], but has not seen wide use within the lidar community. A method to retrieve ice water path and effective radius in cirrus clouds from coincident, space-borne lidar and radar measurements was developed in Delanoe and Hogan [2008], though its results were found to be highly dependent on the microphysical assumptions. Pounder et al. [2012], working in the same group at the University of Reading, outlined a technique to retrieve extinction from three simultaneous observations with different fields of view in a Gauss-Newton iteration using a linearised model of the lidar equation including multiple scattering. This produced high-quality retrievals, though a Twomey-Tikhonov smoothing was added (a term within the cost function penalising the second-derivative

of the retrieved profile). Marchant et al. [2010] presented an original, if limited, linearised retrieval scheme which decomposes the lidar profile into scattering by a basis of representative aerosols (using precomputed scattering properties). This was extended to the retrieval of effective radius in multiwavelength studies via a Kalman filter in Marchant et al. [2012].

A related method known as regularisation (a means to solve Volterra integral equations) has also been applied to the estimation of extinction from Raman lidar profiles to address the underconstrained calculation of the derivative in (2.1.15). The introduction of Veselovskii et al. [2002] provides a good review of early attempts and the methodology. Shcherbakov [2007] demonstrated (using simulated data) that such methods not only return solely positive extinction, but can produce more accurate retrievals. However, that study required both the application of weighted smoothing splines to the observations and a Twomey-Tikhonov smoothing within the retrieval, which generated significant errors near gradients in the extinction. Pornsawad et al. [2012] improved upon that work by removing the smoothing splines and considering only the Raman channel to achieve results of similar quality while proudly proclaiming its believed lack of *a priori* assumptions.

However, such techniques require a set of basis functions to be defined on which to evaluate the integral which artificially impose a structure onto the observations and solutions. Such artificial imposition of smoothing constraints is conceptually unsatisfying. Provided there exists sufficient information on which to base one's assumptions, it is surely preferable to impose the basis by which to smooth solutions through an *a priori* covariance matrix derived from actual data and the physical processes driving the system. In addition, the optimal estimation scheme provides a means to assess the impact of these assumptions through the averaging kernel, such that it is clear where the data dominate the form of the solution. Regularised retrievals do not provide such a balance and, in fact, rarely discuss the choice of basis functions as they are such a fundamental aspect to any practical solution that altering them appears to be a significant task. Further, these authors prefer to evaluate their error budget by Monte Carlo techniques without a thorough consideration of the covariance matrix returned by the algorithm³.

³This may be due to their use of a Twomey-Tikhonov smoothing matrix, which is singular and so cannot be considered the inverse of a covariance matrix.

2.2 Forward model

This thesis proposes an optimal estimation retrieval scheme for aerosol scattering properties from the elastic and nitrogen Raman channels, as described by (1.3.18) and (1.3.19). Estimates of the laser energy, the calibration function, and the number density profile will be required. An Ångström coefficient of unity will be assumed. The background radiation and dark counts will be estimated by averaging over the top 500 bins of the τ_d -corrected profile. The range axis R is assumed known.

Adopting the typical basis of extinction and backscatter, the lidar equations can be written including the background radiation or other noise E_B as,

$$E_i^{\text{el}} = E_L \frac{C_i^{\text{el}}}{R_i^2} \left[\frac{\sigma_R^{\text{el}}}{\mathcal{B}} N_i + \beta_i \right] \exp \left[-2 \int_0^{R_i} \sigma_R^{\text{el}} N(R') + \alpha(R') \, dR' \right] + E_B^{\text{el}} \quad (2.2.1)$$

$$E_i^{\text{ra}} = E_L \frac{C_i^{\text{ra}}}{R_i^2} N_i \exp \left[- \int_0^{R_i} [\sigma_R^{\text{el}} + \sigma_R^{\text{ra}}] N(R') + \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \alpha(R') \, dR' \right] + E_B^{\text{ra}}, \quad (2.2.2)$$

where E is in units of photon counts and the λ argument of α and β has been dropped for brevity. However, this may not be the most suitable basis for these measurements.

As shown by (1.2.9), extinction and backscatter are both functions of N and so will clearly be correlated. This should be identified within \mathbf{S}_a but may not be easily estimated. Further, the use of correlated variables is guaranteed to accentuate degenerate states of the forward model, which can slow retrieval convergence. An alternative variable is the lidar ratio, which is independent of N , for which the forward model would be written,

$$E_i^{\text{el}} = E_L \frac{C_i^{\text{el}}}{R_i^2} \left[\frac{\sigma_R^{\text{el}}}{\mathcal{B}} N_i + \beta_i \right] \exp \left[-2 \int_0^{R_i} \sigma_R^{\text{el}} N(R') + \beta(R') B(R') \, dR' \right] + E_B^{\text{el}} \quad (2.2.3)$$

$$E_i^{\text{ra}} = E_L \frac{C_i^{\text{ra}}}{R_i^2} N(R) \exp \left[- \int_0^{R_i} [\sigma_R^{\text{el}} + \sigma_R^{\text{ra}}] N(R') + \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \beta(R') B(R') \, dR' \right] + E_B^{\text{ra}}. \quad (2.2.4)$$

In either case, the integrals are evaluated by a simple trapezium rule scheme. However, in the lowest range bins the assumptions outlined in Section 1.3.2 are not valid. This should hypothetically be accounted for in the estimation of the calibration function, but likely with an increased error. As such, the retrieved range axis shall begin some distance above the instrument. The extinction will be assumed to be constant through this layer, such that the first bin acts as a boundary term rather than a physically meaningful value.

Then, the retrieved variables must be interpolated onto the instrument's vertical grid. This shall be done after integrating the extinction, though the order should make no difference to the retrieval. Each of the retrieved variables can be cast on to their own vertical grid, but for the lidar ratio to be more meaningful, it is sensible to use the same axis for both variables. An irregularly-spaced grid could be used, but for simplicity, even-spacing will be applied. In theory, provided the grid size is smaller than the features to be resolved, its choice only impacts the computational cost, which is not a great concern at this stage.

Hence, if the variable or superscript r represents the retrieved grid, the aerosol optical thickness is then,

$$\chi_i \equiv \int_0^{R_i} \alpha(R') dR' = \text{spline}_{r \rightarrow R_i} \left[\int_0^r \alpha(r') dr' \right] \equiv \text{spline}_{r \rightarrow R_i} [\chi(r)] \quad (2.2.5)$$

$$\chi(r_l) \equiv \chi_l^r = r_0 \alpha_0^r + \frac{1}{2} \sum_{k=1}^l [\alpha_k^r + \alpha_{k-1}^r] [r_k - r_{k-1}]. \quad (2.2.6)$$

The cubic spline method of Press et al. [1992] is used,

$$\begin{aligned} \text{spline}_{r \rightarrow R_i} [\chi(r)] &= A_{ij} \chi_j^r + B_{ij} \chi_{j+1}^r + C_{ij} (\chi_j^r)'' + D_{ij} (\chi_{j+1}^r)'' \\ A_{ij} &= \frac{r_{j+1} - R_i}{r_{j+1} - r_j} & B_{ij} &= 1 - A_{ij} = \frac{R_i - r_j}{r_{j+1} - r_j} \\ C_{ij} &= \frac{(A_{ij}^3 - A_{ij})(r_{j+1} - r_j)^2}{6} = \frac{(r_{j+1} - R_i)(R_i - r_j)(R_i + r_j - 2r_{j+1})}{6(r_{j+1} - r_j)} \\ D_{ij} &= \frac{(B_{ij}^3 - B_{ij})(r_{j+1} - r_j)^2}{6} = \frac{(R_i - r_{j+1})(R_i - r_j)(R_i + r_{j+1} - 2r_j)}{6(r_{j+1} - r_j)}, \end{aligned} \quad (2.2.7)$$

where $(\chi)''$ indicates the second derivative and r_j is just below R_i .

By requiring that the fit has a continuous first derivative it can be shown that,

$$\mathbf{E}(\chi^r)'' = \mathbf{F}\chi^r, \quad (2.2.8)$$

where,

$$\begin{aligned} E_{(i-1)i} &= \frac{r_i - r_{i-1}}{6} & E_{ii} &= \frac{r_{i+1} - r_{i-1}}{3} & E_{(i+1)i} &= \frac{r_{i+1} - r_i}{6} \\ F_{(i-1)i} &= \frac{1}{r_i - r_{i-1}} & F_{ii} &= -\frac{r_{i+1} - r_{i-1}}{(r_{i+1} - r_i)(r_i - r_{i-1})} & F_{(i+1)i} &= \frac{1}{r_{i+1} - r_i}, \end{aligned}$$

for $1 \leq i \leq l-2$ with l being the number of retrieved levels. This can be solved with standard linear algebra assuming the boundary condition $(\chi_0^r)'' = (\chi_{l-1}^r)'' = 0$.

Finally, a model of the detector's nonlinearity is required. The most widely used and simplest correction is the nonparalyzable detector [Whiteman, 2003; New-som et al., 2009],

$$\varphi^{\text{el,ra}} = \frac{ME^{\text{el,ra}}}{1 + \tau_d \tau_b^{-1} E^{\text{el,ra}}}, \quad (2.2.9)$$

where φ is the measured number of photon counts; τ_d is the dead time; and M is the number of laser shots the profile is averaged over and the photon counts have been divided by the bin width to give a count rate.

Hence, the measurement and state vectors are,

$$\mathbf{y} = \begin{pmatrix} \varphi_0^{\text{el}} \\ \varphi_1^{\text{el}} \\ \vdots \\ \varphi_{m-1}^{\text{el}} \\ \varphi_0^{\text{ra}} \\ \varphi_1^{\text{ra}} \\ \vdots \\ \varphi_{m-1}^{\text{ra}} \end{pmatrix} \quad \text{and} \quad \mathbf{x} = \begin{pmatrix} \beta_0^r \\ \beta_1^r \\ \vdots \\ \beta_{n-1}^r \\ \alpha_0^r \\ \alpha_1^r \\ \vdots \\ \alpha_{n-1}^r \end{pmatrix} \quad \text{or} \quad \mathbf{x} = \begin{pmatrix} \beta_0^r \\ \beta_1^r \\ \vdots \\ \beta_{n-1}^r \\ B_0^r \\ B_1^r \\ \vdots \\ B_{n-1}^r \end{pmatrix}.$$

It must be noted that all elements of \mathbf{x} should be positive. However, there is nothing preventing the algorithm retrieving negative values. Though this is not an issue for the Ansmann scheme, here negative extinction will permit optical thickness to decrease with height. There are two solutions to this, whose merits will be weighed in the following sections:

1. Set a lower limit of zero for the state vector, whereby after evaluating (1.5.10), all negative values are set to zero before being input to the forward model.
2. Retrieve the logarithm of \mathbf{x} . This also has the advantage of compressing the dynamic range of the state from several orders of magnitude to one.

A similar issue exists for the lidar ratio, which is always greater than unity. However, the logarithmic solution will not be considered for any such retrieval as it requires significant manipulation of the forward model to implement and its values cover only an order of magnitude.

One final note must be made of the treatment of measurement error⁴, S_ϵ . The observed photon counts should be Poisson distributed, such that their variance is

⁴Which is assumed uncorrelated.

equal to their mean. In practice, this is widely used to justify approximating the variance of a lidar measurement with the measurement itself. This is not strictly valid, as the measurement is only a single sample of the distribution and a poor estimate of the mean. However, to minimise the storage space and processing time required, lidars sum profiles over several seconds of laser shots online, giving no measure of their variance. At this longer time scale, the variance between successive logged profiles is dominated by atmospheric variations. This may seem a minor point, but the optimal estimation scheme requires an unbiased estimate of the variance. Estimation with the measurement itself causes the retrieval to favour measurements that coincidentally suffer a large, positive noise as these are then presumed to be more precise. This effect is most pronounced at low signal levels near the top of a profile and introduces a high bias into the retrieval. To alleviate this, the variance will be estimated by the application of a sliding window average to the measurement⁵. The exact impact of this will not be explored in detail, though preliminary studies found that even minimal smoothing of the variance vastly reduced residuals⁶.

Summary

The forward model is,

$$\varphi_i^{\text{el,ra}} = \frac{ME_i^{\text{el,ra}}}{1 + \frac{\tau_d}{\tau_b} E_i^{\text{el,ra}}} \quad (2.2.10)$$

$$E_i^{\text{el}} = E_L \frac{C_i^{\text{el}}}{R_i^2} \left[\frac{\sigma_R^{\text{el}}}{\mathcal{B}} N_i + \text{spline}_{r \rightarrow R_i}(\beta^r) \right] \exp[-2(\sigma_R^{\text{el}} \mathcal{N}_i + \chi_i)] + E_B^{\text{el}} \quad (2.2.11)$$

$$E_i^{\text{ra}} = E_L \frac{C_i^{\text{ra}}}{R_i^2} N_i \exp \left[-(\sigma_R^{\text{el}} + \sigma_R^{\text{ra}}) \mathcal{N}_i - \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \chi_i \right] + E_B^{\text{ra}}, \quad (2.2.12)$$

where,

$$\mathcal{N}_i = \int_0^{R_i} N(R') \, dR' \quad (2.2.13)$$

$$\chi_i = \text{spline}_{r \rightarrow R_i}(\chi^r). \quad (2.2.14)$$

The evaluation of \mathcal{N} will be discussed in the next chapter.

⁵Over five bins in all applications to be presented.

⁶Being the difference $\mathbf{y} - \mathbf{F}$ between the measurement and the forward modelled data returned by the retrieval.

The extinction configuration uses,

$$\chi_l^r = r_0 \alpha_0^r + \frac{1}{2} \sum_{k=1}^l [\alpha_k^r + \alpha_{k-1}^r] [r_k - r_{k-1}], \quad (2.2.15)$$

whilst the lidar ratio configuration uses,

$$\chi_l^r = r_0 \beta_0^r B_0^r + \frac{1}{2} \sum_{k=1}^l [\beta_k^r B_k^r + \beta_{k-1}^r B_{k-1}^r] [r_k - r_{k-1}]. \quad (2.2.16)$$

The form of \mathbf{K} is summarised in Section A.1.

2.3 Expected properties and distribution

Arguably the most important component of an optimal estimation scheme is the *a priori*. It constrains which states are deemed to be both physically possible and likely. For example, solutions in this problem should be reasonably smooth, as aerosols should be well-mixed through the PBL. However, gradients in scattering shouldn't be completely excluded as, especially at night, layering is possible. Ideally, the *a priori* would not greatly affect the retrieval, but in practice it will determine which of various degenerate states is preferred as a solution.

2.3.1 Aerosol modelling

The exact composition and optical properties of aerosol are highly variable and climatologies (from which an *a priori* would be derived in other applications) rarely exist. However, some generalised descriptions of representative aerosol types have been explored in the literature. The OPAC (optical properties of aerosol and cloud) model [Hess et al., 1998] defines size distributions, refractive indices, and number densities for a variety of cloud and aerosol particles. This provides the necessary inputs for Mie codes, from which the extinction and backscatter can be readily calculated. Various combinations of these based on expected and observed compositions then produce characteristic aerosol mixtures, such as marine or urban.

For the data to be considered, the continental type should be appropriate, comprising soot with soluble and insoluble aerosols. An ensemble of scattering properties was constructed by randomising the abundance of these components, using the OPAC model values as the mean of a Gaussian distribution with width estimated by a fraction of that mean. Two such estimates (10 % and 20 %) are shown in fig. 2.3 for four cases — a mixture of clouds and spheroidal aerosol, only cloud

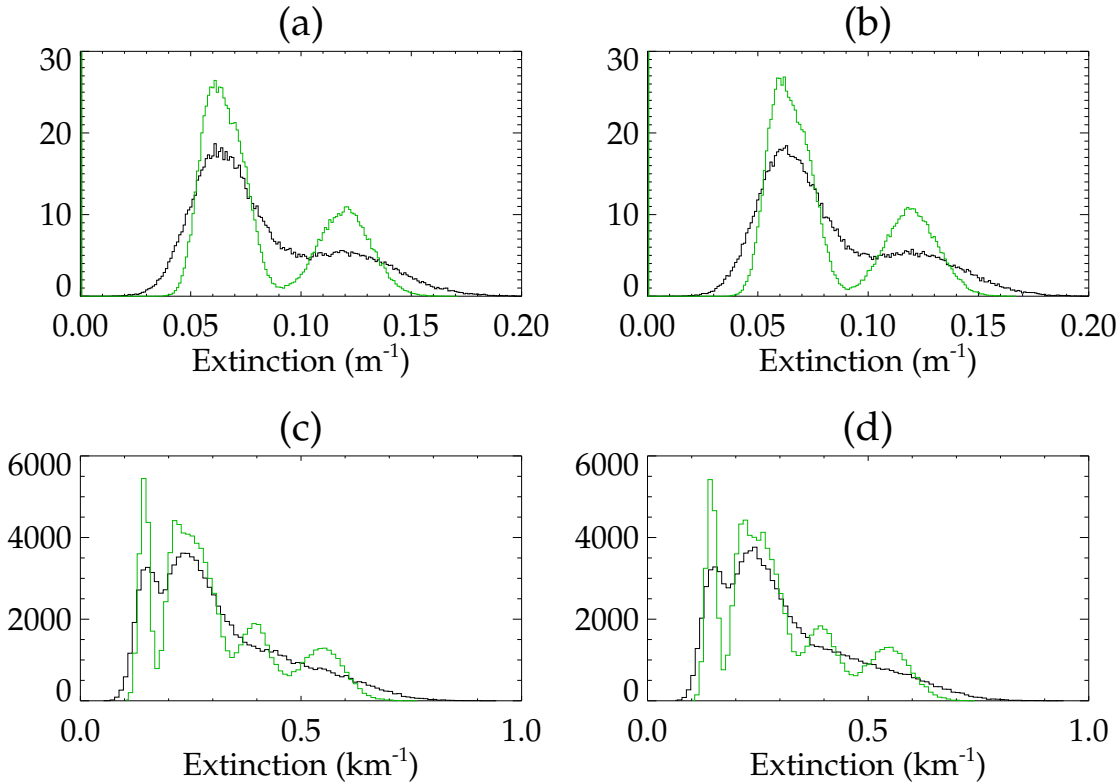


Figure 2.3: The distribution of extinction generated by assuming Gaussian-distributed number densities with width equal to 10 % (green) and 20 % (black) of the mean for (a) Cloud and spheroidal aerosol; (b) Cloud nuclei only; (c) Spheroidal aerosol; (d) Spherical aerosol.

nuclei, only spheroidal aerosol⁷, and only spherical aerosol. The width of the assumed Gaussians, as should be expected, only affects the width of the features of the distribution and not its general shape. The impact of aspect ratio is negligible.

The 10 % distribution for either cloud and aerosol or only aerosol are then used as a basis for the *a priori*. The *a priori* state and variance represent a Gaussian distribution for each state vector element with mean $(\mathbf{x}_a)_i$ and variance $(\mathbf{S}_a)_{ii}$, values of which can be estimated by fitting to the above distributions. Rough fits have been performed by eye and are plotted in figure 2.4 in blue. The red curves show similar fits but for when the log of the state vector is retrieved instead, which clearly better models the predicted distributions of extinction and backscatter. Though the lidar ratio is theoretically a better description of the state, its distribution is not symmetric and so not necessarily well-suited to optimal estimation. In an attempt

⁷The aspect ratio was assumed to be uniformly distributed between 21 discrete values on the range 0.335-2.986 considered in the T-matrix code [Dubovik et al., 2006].

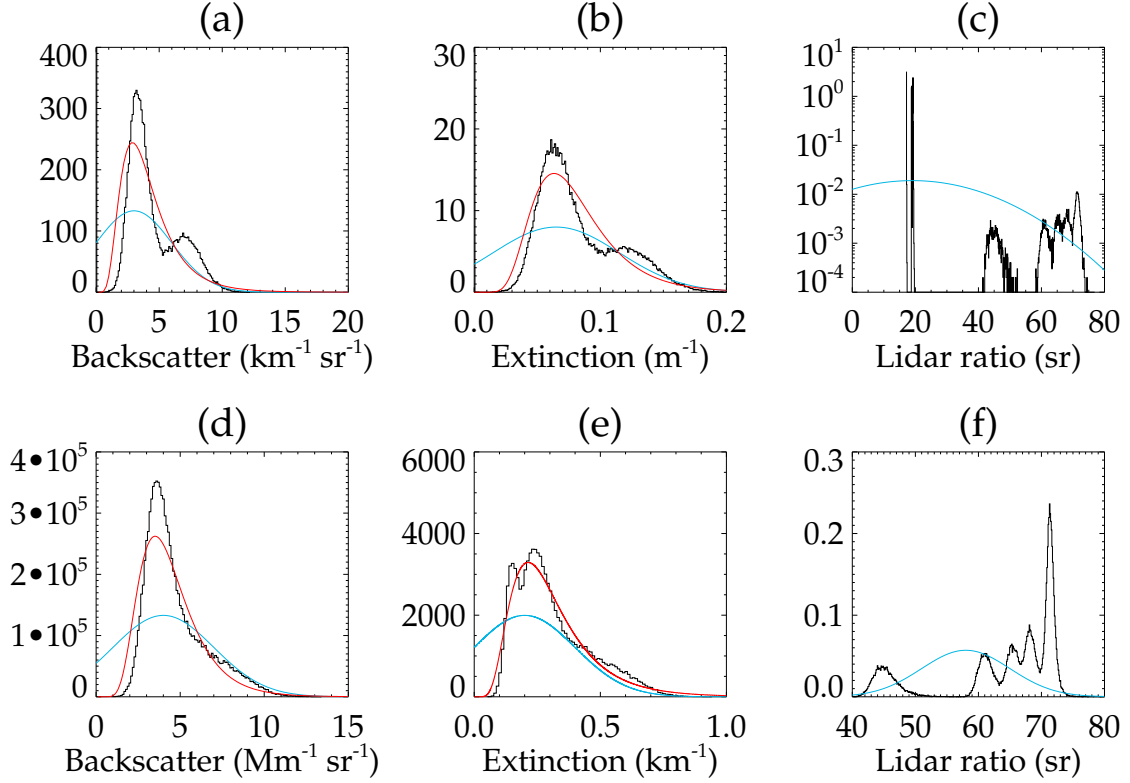


Figure 2.4: *A priori* aerosol distributions, with the black curve reproduced from fig. 2.3. The blue curve represents the linear x model and the red $\ln x$. Plots (a-c) are for a combination of cloud and aerosol whilst (d-f) show spheroidal aerosol. Parameter values summarised in table 2.1.

to compensate for this, relatively broad *a priori* distributions have been selected.

The OPAC model states that the density of non-dust aerosols decreases exponentially with a scale height of 2 km. To represent this, x_a will decrease from its calculated value similarly. Further, there will almost certainly be some vertical correlation of the measurements due to vertical motion of the particles⁸. The simple model of a Markov process proposed by (2.83) of Rodgers [2000] shall be used with correlations decaying exponentially with separation,

$$(\mathbf{S}_a)_{ij} = \sqrt{(\mathbf{S}_a)_{ii}(\mathbf{S}_a)_{jj}} \exp\left(-\frac{|r_i - r_j|}{H}\right), \quad (2.3.1)$$

where H is some scale height. The first bin is not correlated with any other due to its greater extent, as mentioned earlier.

⁸Subsequent profiles should also have some correlation in time, but the scheme to be developed will consider profiles individually as the matrix inversions already require significant processing.

Table 2.1: *A priori* aerosol distributions for α/β retrieval, as shown in fig. 2.4.

		Cloud			Aerosol		
		β	α	B	β	α	B
Linear	Mean	3×10^{-3}	6.5×10^{-2}	19	4×10^{-6}	2×10^{-4}	58
	SD	3×10^{-3}	5×10^{-2}	21	3×10^{-6}	2×10^{-4}	7
Log	Mean	-5.6	-2.6		-12.4	-8.2	
	SD	0.5	0.4		1.0	0.5	

The scale height can be estimated by investigating the covariance of some measure of aerosol scattering. A convenient option is backscatter sondes [NDACC [2000]], which measure the light backscattered from a xenon flashlamp approximately every 30 m during a balloon ascent [Rosen and Kjome, 1991; Rosen et al., 2000]. Profiles over ten years of observations at three sites have been used to produce a correlation matrix of backscatter ratio with height (fig. 2.5). Its rows decay roughly exponentially with height. The plot presents a number of box-like features produced by layers of unusually large aerosol concentration observed on a single launch⁹, dominating the average β at that height. However, this doesn't greatly impact the result of a least squares fit of (2.3.1) to each row of the matrix (right of fig. 2.5), showing that correlations are of order kilometres in the free troposphere.

However, this will not necessarily apply within the PBL, which is only weakly coupled to the free troposphere [Oke, 1987]. Several studies of aerosol vertical distribution within the PBL have been performed with tethered balloons, though the data could not be readily accessed. These generally find that aerosol concentrations are constant with height [figs. 10–12, 4, and 2 of Ewell et al., 1989; Greenberg et al., 2009; Ferrero et al., 2010, respectively], but occasionally observe fine structure [fig. 6 of Ferrero et al., 2011]. Further, a myriad of literature covers observations with lidar of aerosol layers tens to hundreds of metres thick [e.g. Di Girolamo et al., 1999; Ansmann et al., 2010; Dacre et al., 2011] or variations within lofted aerosol features [e.g. Althausen et al., 2000; Huang et al., 2010].

Clearly, the *a priori* covariance matrix should represent both the general tendency for aerosols to be well-mixed throughout the troposphere and the fine scale structure that occasionally occurs. Further, the average position of the top of the PBL should be represented by a significant decrease in correlation between areas above and below it. However, at the moment, there is insufficient information to

⁹Sondes were almost exclusively launched when such an event was expected.

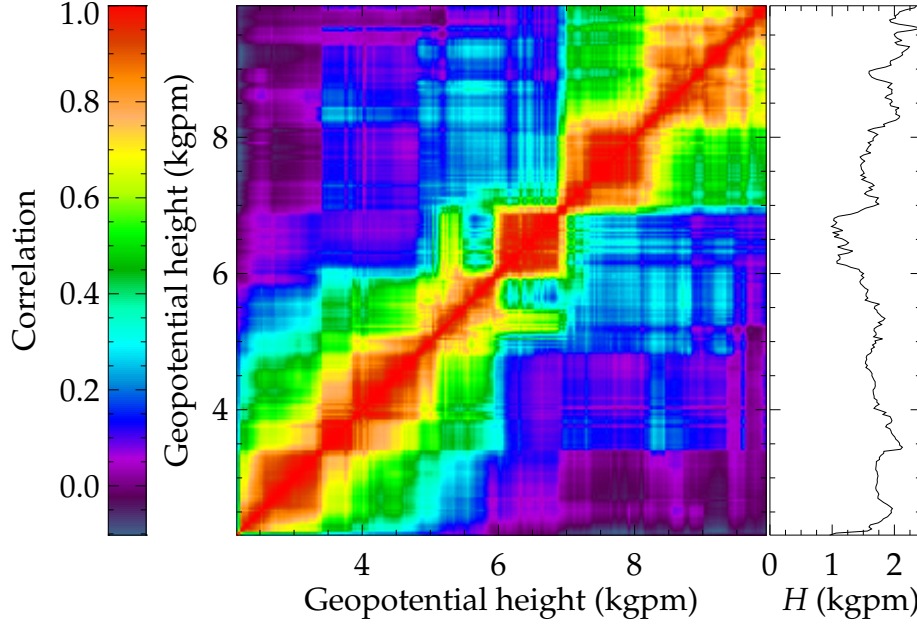


Figure 2.5: The observed vertical correlation of backscatter. **Left** — Autocorrelation of backscatter with height, derived from 198 backscattersonde profiles collected between 1989 and 2000 at Laramie, WY, USA; Lauder, New Zealand; and Thule, Greenland. **Right** — Least-squares fit of (2.3.1) to each row of that matrix.

represent these effects with any degree of certainty. As such, it is preferred to use a conservative estimate of $H = 100$ m, which will in most cases not make the best use of the available information but has a lower risk of overconstraining the solution.

2.3.2 Simulated data

Simulated data will be used to assess the various proposed configurations of the retrieval. Under such controlled conditions, the strengths and weaknesses of the retrieval can be broadly assessed, within the limits of what can be reasonably simulated. To this end, six representative profiles have been created by evaluating equations (1.3.14) and (1.3.15). To summarise the features of each profile (outlining the contents of table 2.2),

1. The basic model profile with constant lidar ratio and idealised PBL structure. The system is reasonably well-aligned and the observations are made at night with purely nonparalyzable detectors.
2. As case 1, but with lower laser energy and higher background count (approximating dusk) to explore the impact of noise on the retrieval.

3. As case 1, but with a different extinction profile. The PBL is taller with greater extinction.
4. As case 1, but with a larger lidar ratio. Further, a layer of aerosol with distinct lidar ratio is added at 800 m.
5. Using the same extinction profile and lidar ratio as profile 1, but for a slight misalignment of the instrument. In addition, the detectors are paralyzable in the manner outlined in (1.3.24).
6. Observation of a cloud. The PBL is lower and less optically thick than in profile 1 with a liquid water cloud situated 1.1 km above it's top. The slightly misaligned instrument of case 5 is retained, but observed with nonparalyzable detectors.

All profiles used the same standard atmosphere [NOAA [1976]], but with different, randomised surface pressure and temperature (table B.3). The effective area of the instrument was determined using (4.2.22).

The idealised PBL extinction profile of Steyn et al. [1999] is adapted to account for decreasing aerosol concentrations in the free troposphere by multiplying by an exponential decay (with scale height H_{FT}) above the PBL,

$$\alpha(z) = \begin{cases} \frac{1}{2}(\alpha_{BL} + \alpha_{FT}) + \frac{1}{2}(\alpha_{BL} - \alpha_{FT}) \operatorname{erf}\left(\frac{z_{BL}-z}{\sigma_{BL}}\right) & z \leq z_{BL}, \\ \left[\frac{1}{2}(\alpha_{BL} + \alpha_{FT}) + \frac{1}{2}(\alpha_{BL} - \alpha_{FT}) \operatorname{erf}\left(\frac{z_{BL}-z}{\sigma_{BL}}\right) \right] \exp\left(\frac{z_{BL}-z}{H_{FT}}\right) & z > z_{BL}, \end{cases} \quad (2.3.2)$$

where α_{BL} and α_{FT} are the extinction in the PBL and free troposphere, respectively; z_{BL} is the PBL height; and σ_{BL} is the width of the entrainment layer.

Table 2.2: Details of simulated profiles. For each element of the lidar equation that is varied, the version used for that profile is denoted (details of which can be found in tables 1.1, B.2, and B.3).

Profile #	E_L	$C(R)$	$\alpha(R)$	$B(R)$	$N(R)$	τ_d	E_B
1	1	1	1	1	1	1	1
2	2	1	1	1	2	1	2
3	1	1	2	1	3	1	1
4	1	1	1	2	4	1	1
5	1	2	1	1	5	2	1
6	1	2	3	3	6	1	1

The aerosol and cloud layer are modelled by a Gaussian [Biavati, G., personal correspondence],

$$\alpha_{\text{feature}}(z) = \frac{\alpha_{\text{feature}}}{\sigma_{\text{feature}} \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{z - z_{\text{feature}}}{\sigma_{\text{feature}}} \right)^2 \right], \quad (2.3.3)$$

with a constant lidar ratio for this component.

Once simulated, Poisson noise is added to the profiles. Though this should be done by sampling a Poisson distribution for each laser shot and taking the total, an equivalent and much faster method is to sample only once a distribution with mean equal to the modelled profile times the number of laser shots.

2.4 Configuration assessment

There are a number of possible configurations for the retrieval and its *a priori* to investigate:

- In addition to backscatter, should it retrieve extinction or the lidar ratio?
- If retrieving extinction, should its correlation with backscatter be estimated?
- Should the log of the state vector be retrieved instead?
- Should the *a priori* be based on aerosol or cloud particles?
- Is the assumed scale height for vertical correlation important?

A simple first example is shown in figure 2.6. It retrieves α and β linearly while neglecting their correlation using the aerosol *a priori*¹⁰ and $H = 100$ m with a lower limit of zero applied to the state vector¹¹. Retrieval from the first four cases of simulated data are broadly reasonable with the backscatter reproducing the simulation accurately. Case 2, as expected, exhibits greater noise than case 1. The underestimation of the dead time correction in case 5 produces a significant underestimation of backscatter within the PBL in a solution which failed to converge. Case 6 is very poorly managed, fitting only half of the data available before the laser is completely attenuated. It also fails between the PBL and the cloud, though performs within the PBL as well as the other cases. Extinction is less well retrieved

¹⁰See table 2.1.

¹¹As will be done for all linear configurations herein.

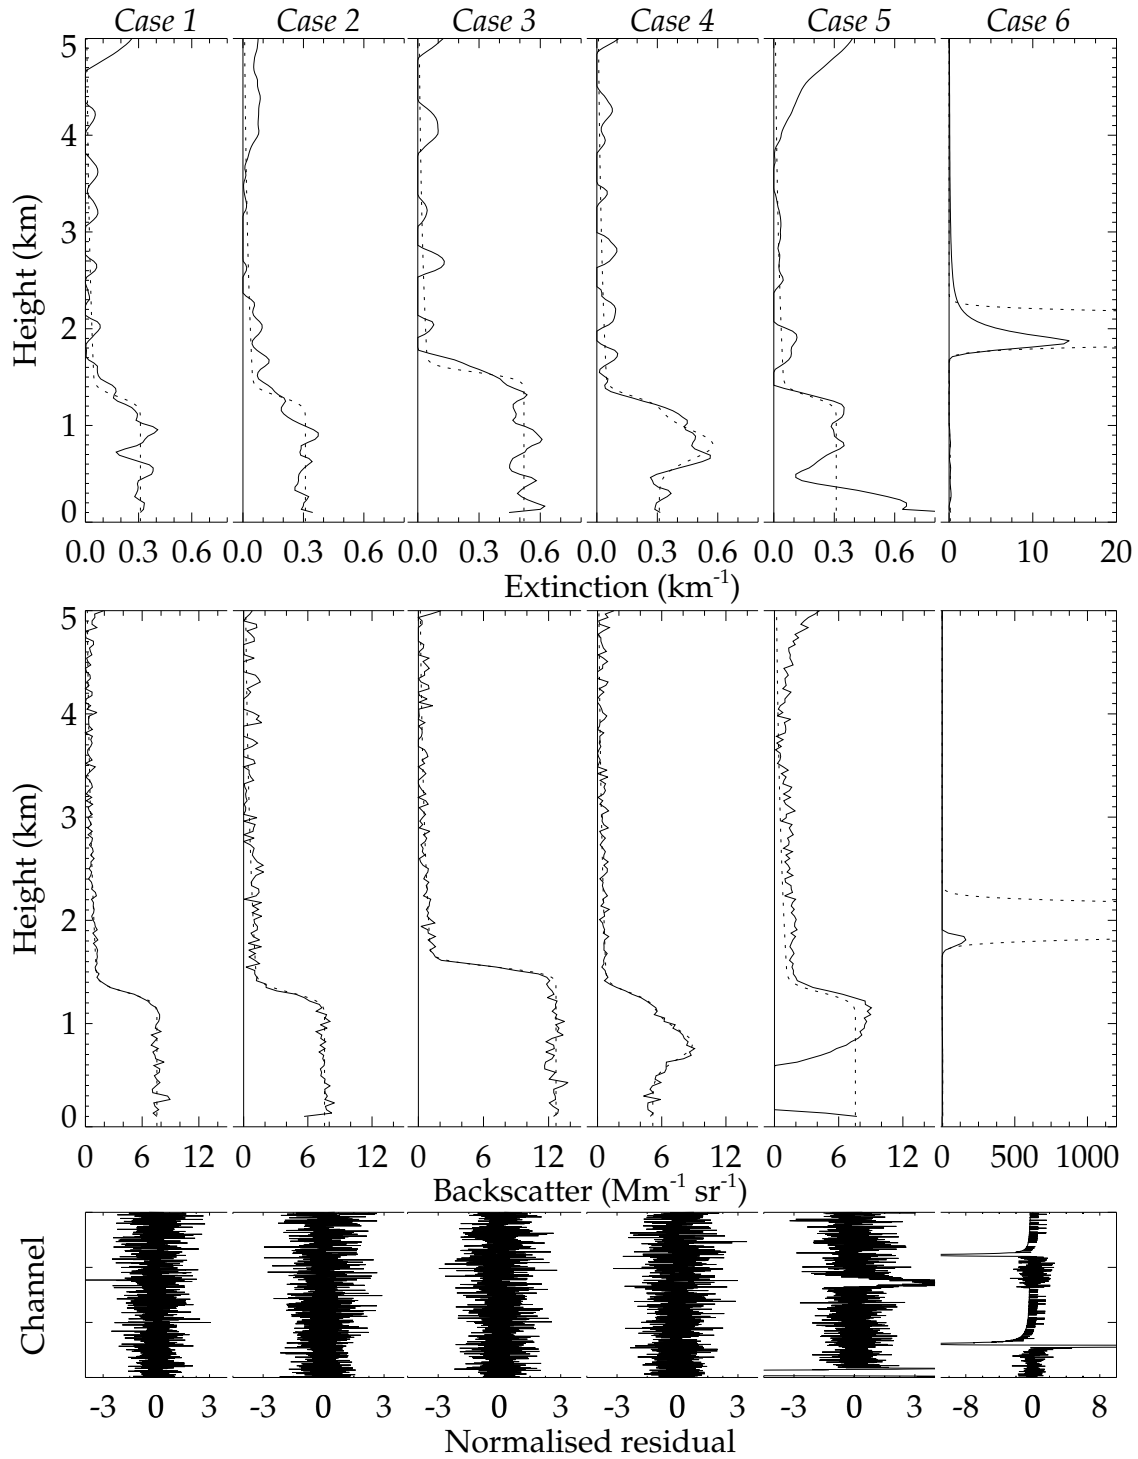


Figure 2.6: A first application of the retrieval, showing profiles of α and β (top and centre) from simulated cases 1–6 (left to right), where their correlation has been neglected and using an aerosol-inspired *a priori*. The dotted line shows the true value. Residuals (bottom) are normalised by S_ϵ and cover the entire measurement vector such that the bottom half are elastic channels and the top half Raman.

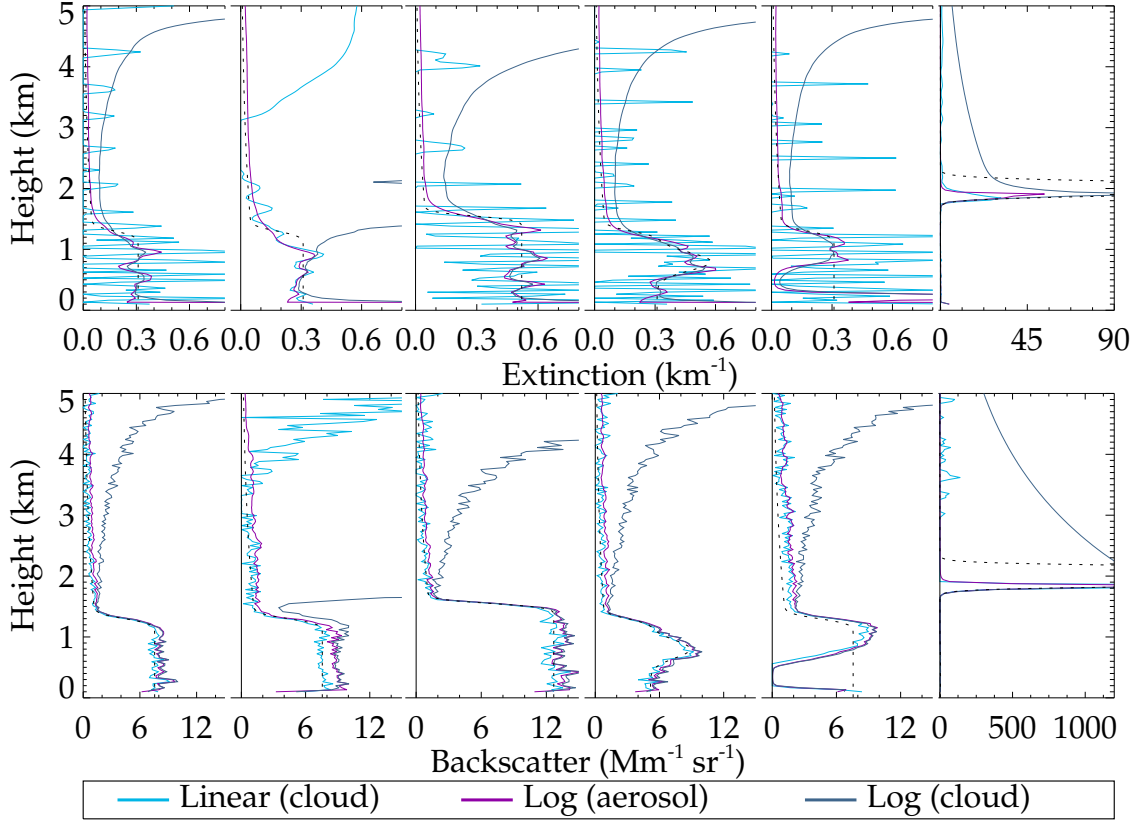


Figure 2.7: Attempts to improve the retrieval of cloudy data by applying different *a priori* derived from cloud or aerosol distributions and/or retrieving the logarithm of x . The dotted line shows the true value.

in all cases. The periodic structure retrieved is seeming coincidental, being independent of the retrieval's parameters. The residuals demonstrate that the forward modelled solutions are unbiased estimates of the measurement, except where the detector saturates in case 5 and within the cloud in case 6.

In an attempt to better fit the cloud, the *a priori* could instead reflect cloud properties. Alternatively or additionally, the logarithm of the state vector could be retrieved to compress the dynamic range and encourage larger values. As shown in figure 2.7, when using either alteration alone, a more appropriate fit to the cloud is achieved¹². Both logarithmic retrievals produce a smoother extinction profile above the PBL than the linear configurations. The increased *a priori* variance favours spikes in the linear retrieval of α , with the drift in case 2 corresponding to the retrieval reverting to the *a priori* state. It can be brought in line with the other solutions by using the aerosol x_a . The use of both corrections simultaneously is

¹²The simulated signal dies away at 1.9 km so the sharp peak is reasonable.

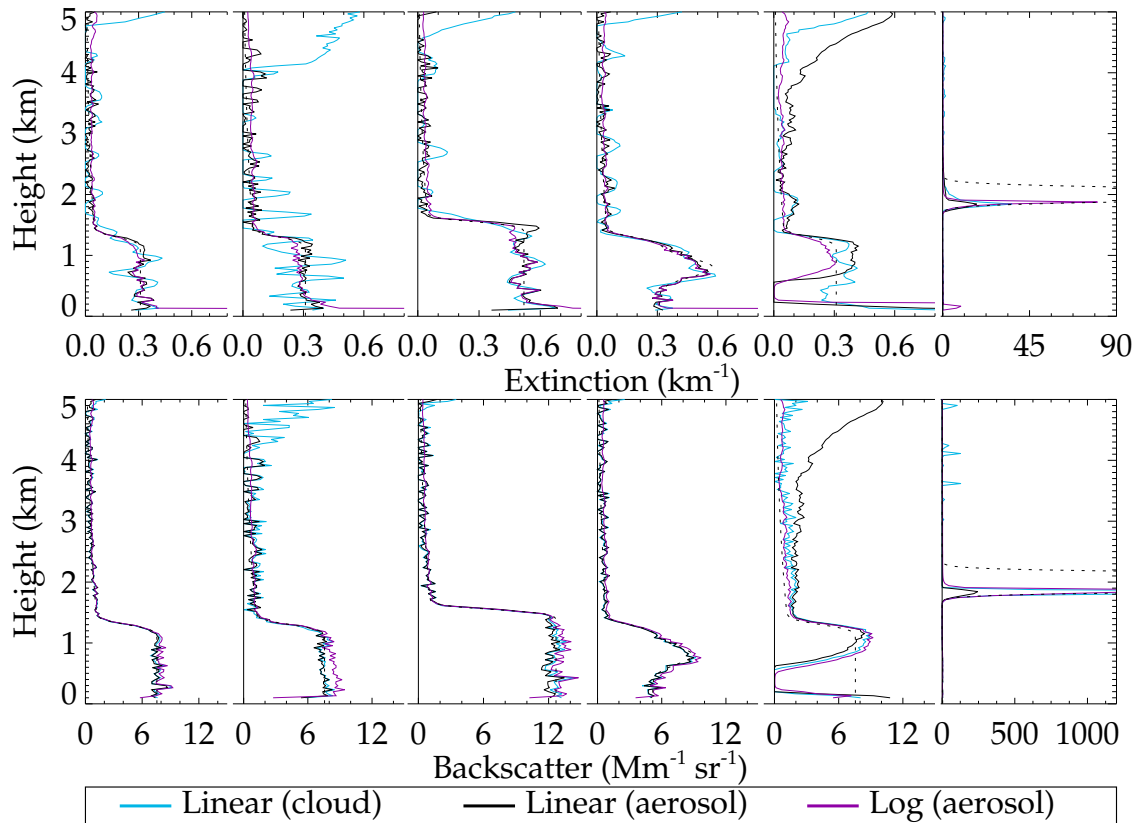


Figure 2.8: As fig. 2.7, but assuming a correlation between extinction and backscatter.

simply disastrous, never able to converge to a solution and tending towards excessive scattering with height.

Generally, it appears that the retrieval does not perform well in the presence of clouds. A greater *a priori* variance allows the retrieval to fit larger values and see further into the cloud, but these are still not large enough to properly fit the data. Using vertical correlation in the *a priori*, the retrieval attempts to compensate by minimising values beneath the cloud (possibly to increase the signal gradient at cloud base). This hinders convergence and produces a practically useless extinction profile. As such, aerosol retrievals shall guide further algorithm development with performance in the presence of clouds being a secondary concern.

It has been noted that this formulation of the problem is unphysical. A 95% intercorrelation between the extinction and backscatter appears to be reasonable from the ensemble of OPAC-derived aerosol scattering properties. Though the exact value assumed appears to have minimal impact on the retrieval, a more rigorous estimation of this value would be desirable. The intercorrelation must decay

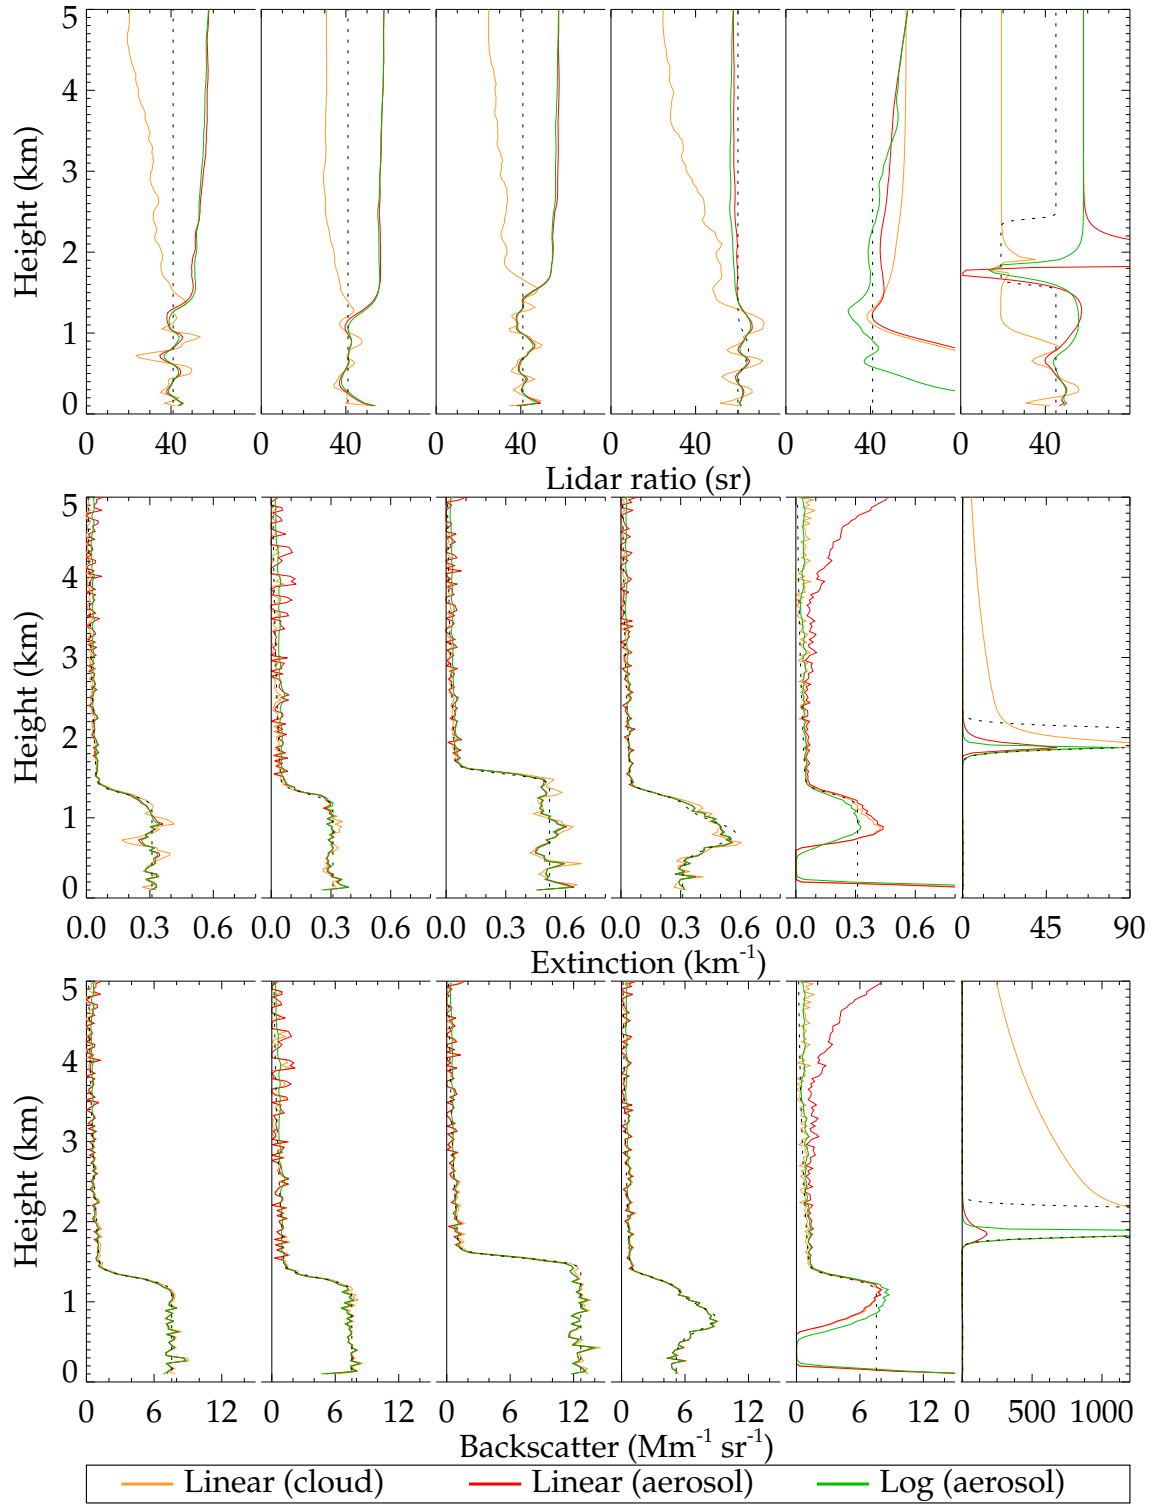


Figure 2.9: Similar to fig. 2.7, but retrieving the lidar ratio instead of extinction.

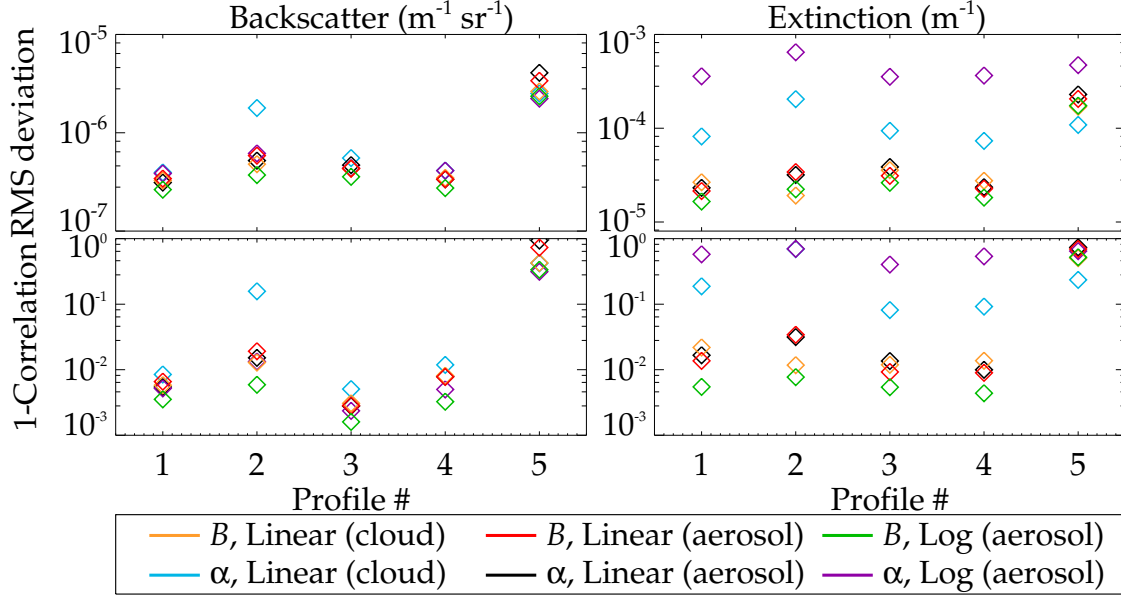


Figure 2.10: Deviation of retrieved from simulated profiles. **Top:** RMSD of the retrieved backscatter (left) and extinction (right) from the simulated profile for the configurations presented in figs. 2.8 and 2.9. **Bottom:** One minus the correlation coefficient of the same.

over the same scale height as vertical correlations or the covariance matrix will be inconsistent (in that it would not be positive definite). Doing so produces figure 2.8, showing the retrieval becomes more consistent with the truth. The logarithmic retrieval continues to demonstrate the minimal response to noise, though tends to bias high¹³. The linear solution is basically a noisier version of that in fig. 2.6.

An alternative formulation would be to retrieve the lidar ratio. As shown in figure 2.9, the resulting profiles are the most suitable yet, being smoother at all levels and no longer presenting a high bias. The three configurations shown perform broadly identically, with each lidar ratio profile tending towards its *a priori* above the PBL. Whether this loss of sensitivity is due to the small magnitude of α above the PBL or is a matter of noise is not yet clear.

The suitability of the various configurations can be quantitatively assessed by the RMSD of the retrieved profile from the true profile or by their correlation. Presented in fig. 2.10, it is clear that the retrieval of the lidar ratio is superior to extinction, with the logarithmic retrieval of backscatter best suiting the former method. If extinction must be retrieved, a linear retrieval is best. In either case,

¹³This bias is small enough to not be reflected in the retrieval cost such that, like every other retrieval of cases 1–4 that doesn’t diverge with increasing height, it has a cost of less than 1.1.

the aerosol *a priori* is superior. Note that the quality of all retrievals is significantly degraded where the dead time is improperly corrected.

The question of correlation scale is more difficult to address. Provided H is greater than the size of a measurement bin, the retrieved profile is reasonably independent of it (fig. B.2). Small scales slightly underestimate values in the PBL and large scales similarly overestimate in the free troposphere, such that 100 m appears to be a reasonable compromise. However, the cases considered thus far do not challenge the resolution of the retrieval.

2.5 Sensitivity

Considering the two remaining configurations for the retrieval, figure 2.11 shows how each responds to noise-free data for two different correlation scales. They retrieve virtually identical backscatter profiles, but the lidar ratio configuration loses sensitivity above the PBL, tending towards its *a priori* value. The extinction configuration maintains sensitivity throughout the profile. However, it is more dependent on H . A large value permits bins in the free troposphere to impact those within the PBL to produce a peak/trough pair about the entrainment layer. Using a value more similar to the width of the entrainment layer minimises this.

By plotting on a log scale, it is evident that the quality of the lidar ratio retrieval is due to the small magnitude of the state above the PBL, such that any discrepancies there are negligible when compared to those in the PBL. Indeed, though the $H = 100$ m extinction configuration is consistent with the truth through almost the entire range considered, it has a higher cost than the other configuration due to the differences in the entrainment layer.

The α mode is also more sensitive to measurement noise. Figure 2.12 shows retrievals for six versions of case 1 with E_L doubled each time. The lidar ratio configuration clearly improves with each increase. Its solutions become smoother and retrieve the lidar ratio to greater heights, which indicates that the loss of sensitivity is a function of SNR. Retrievals in the extinction configuration do improve, but not as markedly. That mode appears to be more accurate but less precise.

This can be further explored with a further set of simulated profiles. Presented in figure 2.13 are:

7. Similar to case 1 but with the addition of ‘aerosol layers’ of the form,

$$\alpha_{\text{feature}} \exp \left[-\ln 2 \left(\frac{z - z_{\text{feature}}}{\sigma_{\text{feature}}} \right)^2 \right] \cos^2 \left[\frac{2\pi(z - z_{\text{feature}})}{\sigma_{\text{feature}}} \right], \quad (2.5.1)$$

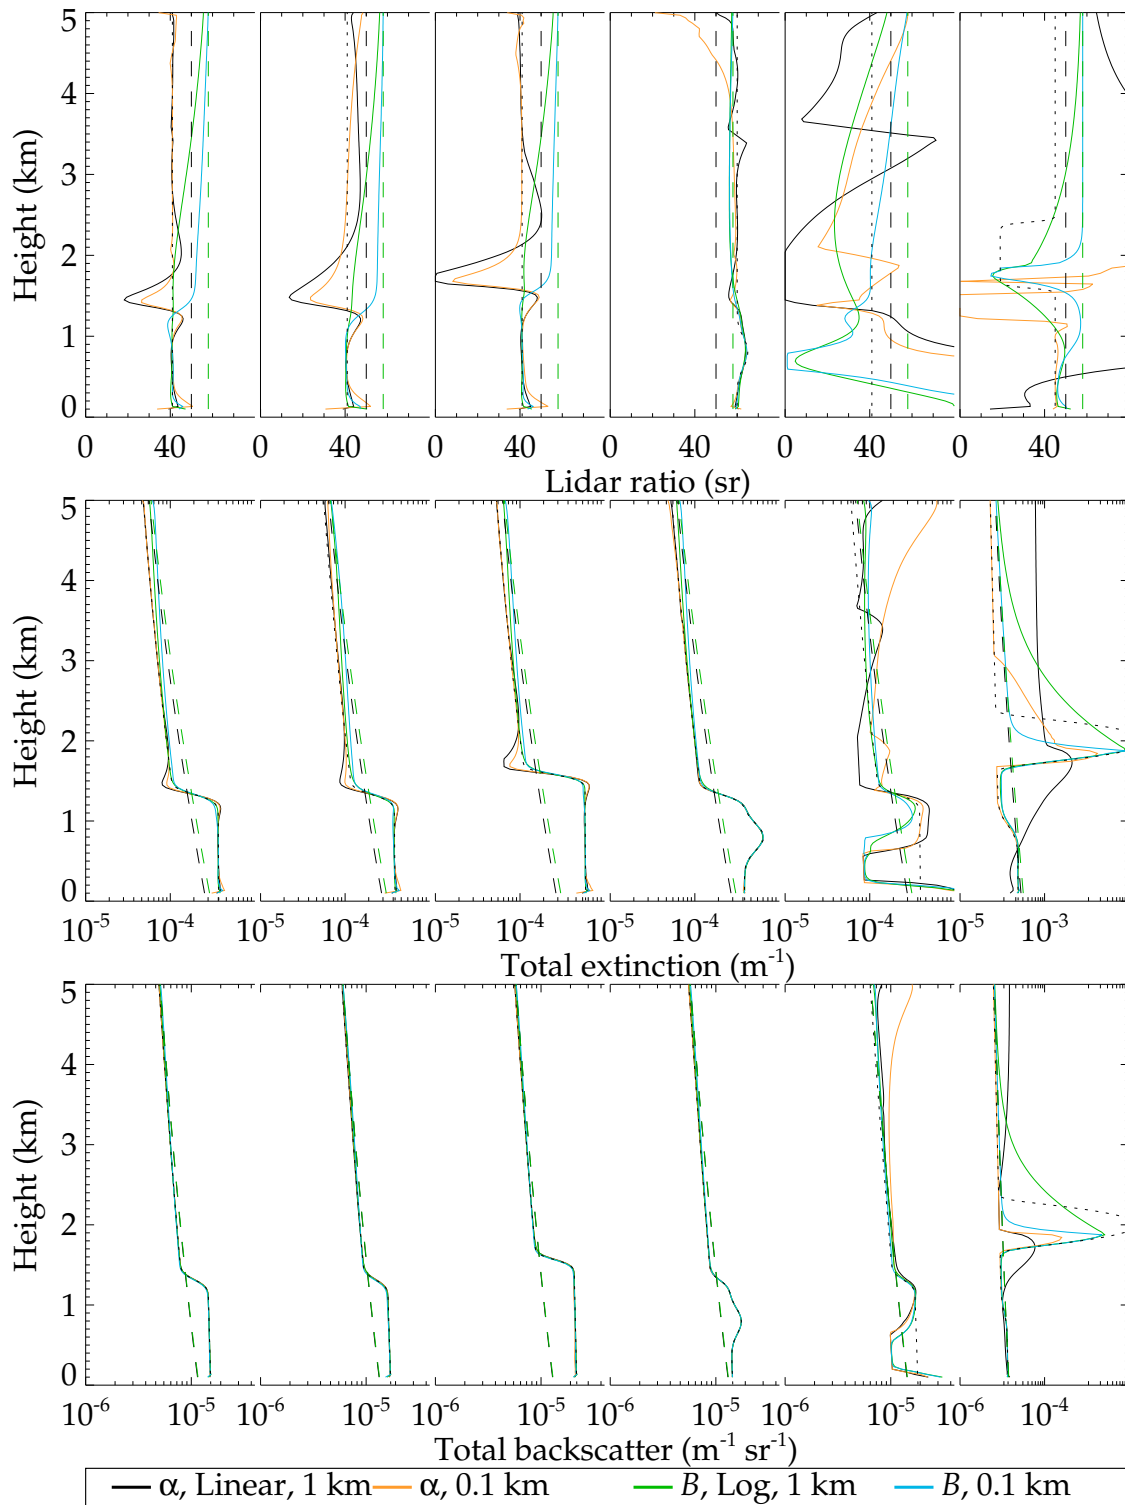


Figure 2.11: Performance of the preferred retrieval configurations with noise-free data for different correlation scale heights. Molecular scattering has been added to extinction and backscatter such that they can be sensibly plotted on a log scale to highlight performance above the PBL. *A priori* profiles plotted as dashed lines.




Figure 2.12: Change in the retrieval of case 1 as the integration time is increased from 5 to 160 minutes in factors of two (left to right) for $H = 100\text{m}$ in the two preferred configurations.

where $\sigma_{\text{feature}} = 300$ m.

8. As case 7, but for $\sigma_{\text{feature}} = 200$ m.
9. As case 7, but for $\sigma_{\text{feature}} = 100$ m.
10. Similar to case 1 with the addition of three distinct and overlapping aerosol layers of the form of (2.3.3).
11. A basic model profile, but with optical thickness 0.12 rather than the 0.50 of case 1.
12. A basic model profile with optical thickness 0.04.

The two configurations behave as before with the lidar ratio mode being smoother than the α mode but losing sensitivity above the PBL. The ‘layers’ of cases 7 and 8 are correctly positioned, if somewhat underestimated in magnitude. In case 9 the layers are not resolved (see fig. 2.14 for detail). The reason for this is simple — the features occupied only one retrieval bin. Doubling the resolution gives equivalent performance to the previous two cases, but with slightly increased noise and significantly increased processing time¹⁴. A less-desirable aspect of assuming a vertical correlation is clear in case 8, where structures appear echoed beneath the actual layers in the backscatter.

Both configurations give a respectable fit to the extinction profile, though they do increasingly underestimate the magnitude of the peaks as they become narrower. This decreased sensitivity is clear from a consideration of the averaging kernels (fig. 2.15). Though the backscatter kernels are virtually delta functions in the PBL, the extinction kernels have a width of ~ 300 m for both configurations. The kernels also illustrate the loss of sensitivity in the lidar ratio configuration above the PBL, with even the magnitude of the backscatter kernels decreasing significantly. In cases 11 and 12, with reduced SNR, the sensitivity is also lower. The extinction configuration maintains sensitivity throughout the profile, though its extinction kernels are strangely skewed about their centre¹⁵. Overall, these indicate that the improved noise performance of the lidar ratio mode is due to a greater reliance on the *a priori*.

¹⁴Fig. B.3 demonstrates that changing the resolution makes virtually no difference to the retrieval of cases 1–6. The selected value is simply a trade-off between processing time and the minimum resolvable scale.

¹⁵This may derive from E^{ra} measuring the integral of α , such that bins beneath a level contribute greater information content, but it is by no means clear.

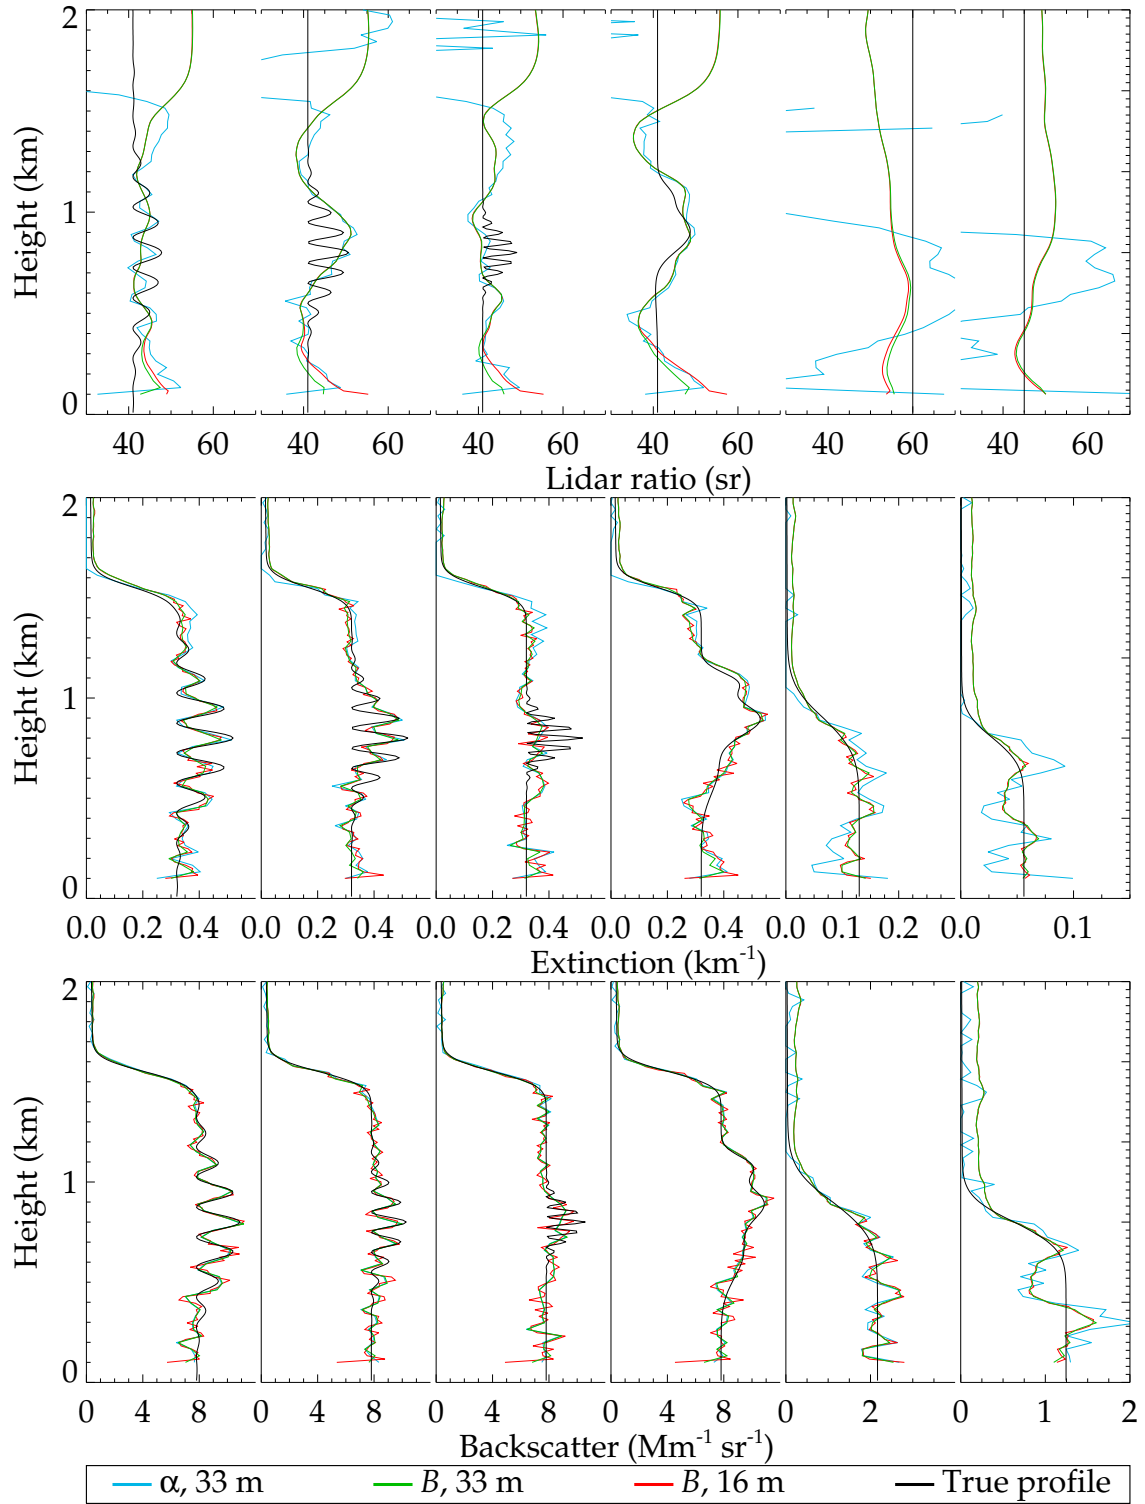
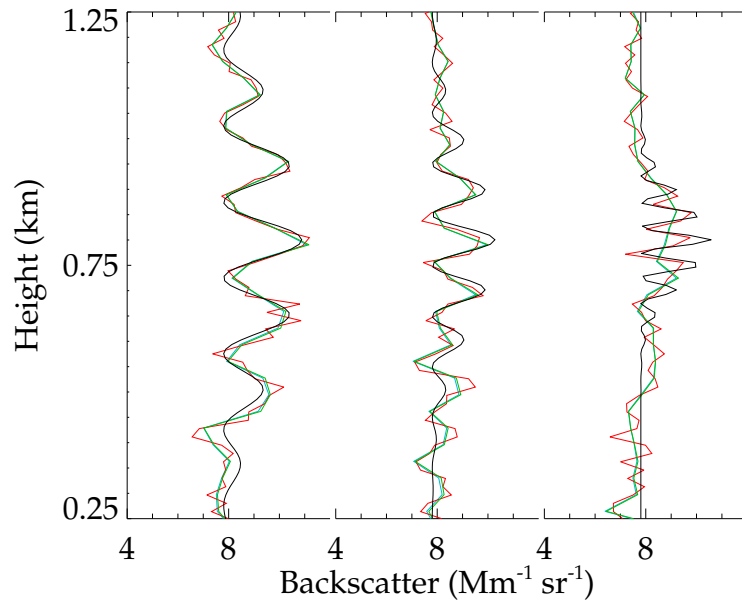


Figure 2.13: Sensitivity to the fine-scale fluctuations of cases 7–12 (left to right), considering the impact of increasing the resolution of the lidar ratio retrieval. The true profile is plotted in solid black for clarity.

Figure 2.14: A closer examination of cases 1–3 of figure 2.13, highlighting how the coarse retrievals fail to capture the smallest features and that all resolutions tend to smooth the magnitude of the peaks.



Returning to figure 2.13, the extinction configuration is just as successful as the lidar ratio at resolving layers, with both smoothing features in the β profile as they become smaller than the retrieved resolution. This implies that the backscatter should be retrieved at the instrument's native resolution and the lidar ratio on a coarser axes. That possibility will not be explored further as the required alterations to the computer codes are significant and the result has been found to be independent of the retrieved axes in preliminary work. This alteration would merely serve to make more efficient use of the computing resources available, which is not deemed a significant concern.

Cases 11 and 12 present more of a challenge with lower levels of signal relative to noise. They are still broadly fitting the true profile, but with more obvious deviations that demonstrate the levels of error in the retrieval and how the two different configurations express it differently.

2.6 Errors

As outlined in Section 1.5.2, optimal estimation also returns an error covariance matrix. Having selected appropriate configurations for the retrieval, it now remains to investigate if that error ensures it is consistent with the simulated profile.

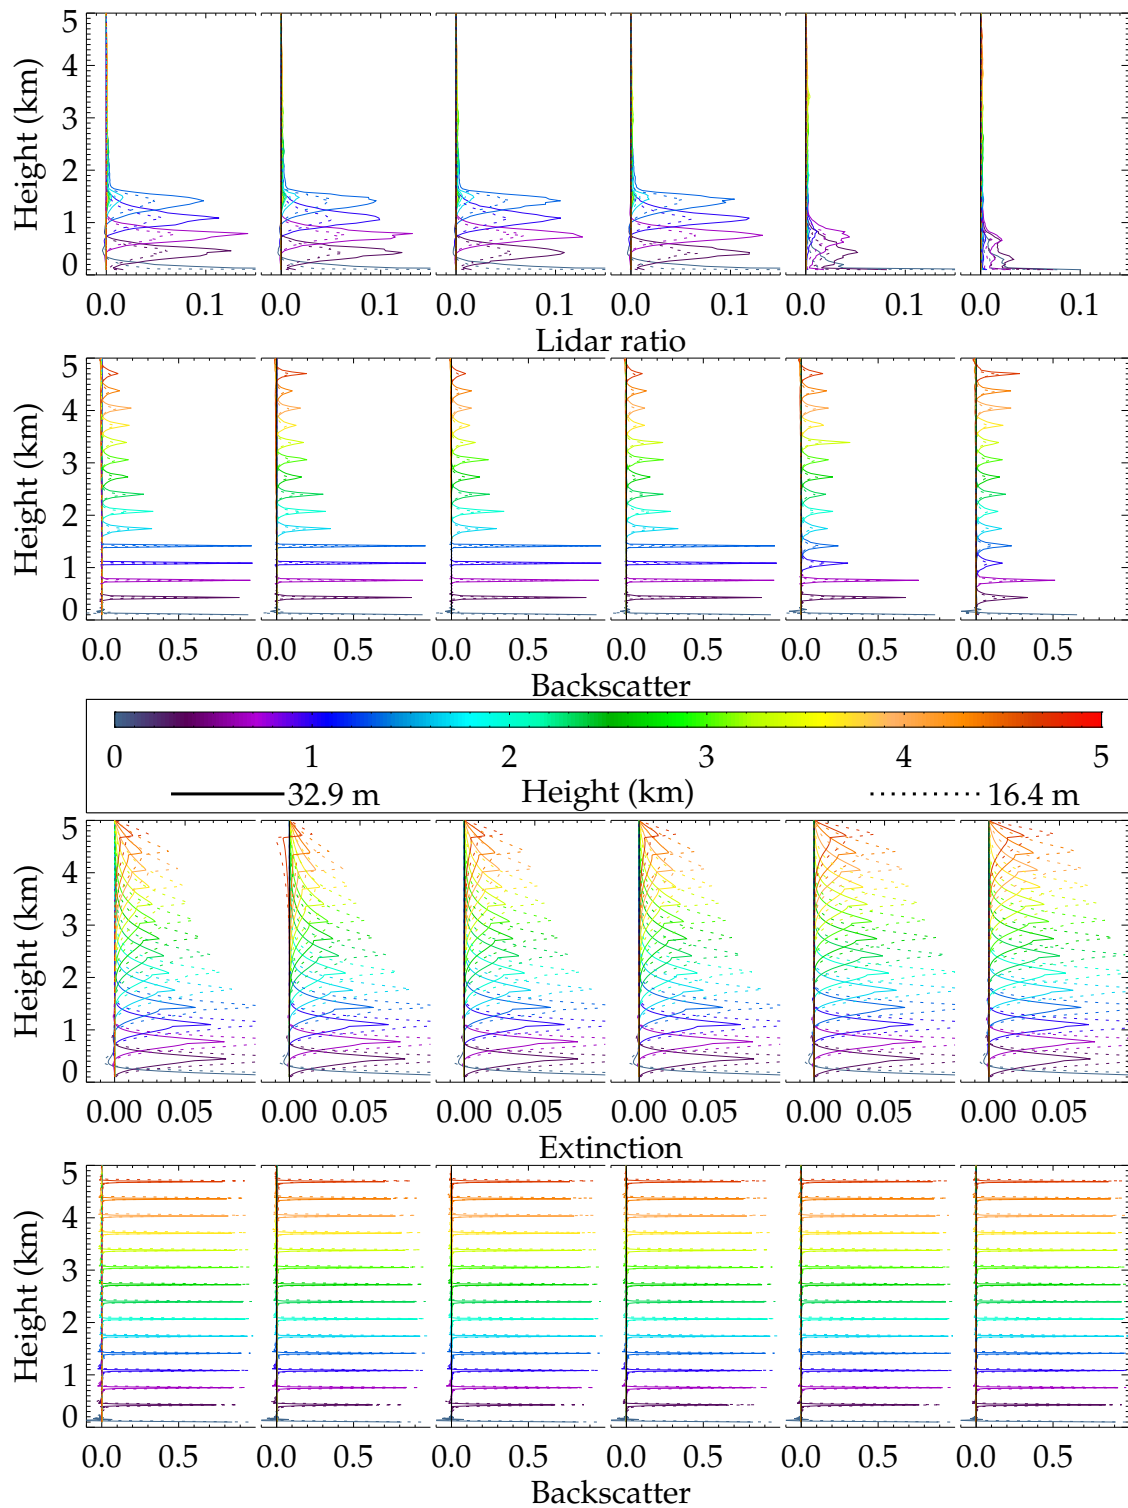


Figure 2.15: Averaging kernels for fig. 2.13. **Above colour bar:** Lidar ratio retrieval. **Below:** Extinction retrieval.

2.6.1 Retrieval error

An example covariance matrix is shown in figure 2.16 for the lidar ratio configuration. The reversion to the *a priori* solution is very clear in the lidar ratio error, plot (b), whilst backscatter shows a more significant improvement over the *a priori*. The form of the covariance matrix has also changed significantly from its *a priori*. Autocorrelation in the backscatter has decreased with regions near the surface being only 10 % correlated to adjacent bins. The extinction matrix, plot (e), shows a more obvious transition from successful retrieval to *a priori*. Near the surface, adjacent bins are correlated to $\sim 60\%$ whilst other bins within the range of the averaging kernel are slightly anticorrelated. This reverts to the *a priori* matrix as the information content diminishes at the top of the PBL. The intercorrelation of the variables has also changed form, with points above a level being anticorrelated and those below positively correlated. It will need to be investigated if real data give a similar form. Equivalent plots for the extinction configuration are presented in figure B.4, highlighting how sensitivity persists throughout the profile and vertical correlations decrease with height.

Plotting this error gives figures 2.17 and 2.18. The lidar ratio retrieval appears to be underestimating its error as it loses sensitivity, but as that error is merely the *a priori* variance, it only indicates that the *a priori* is overly constrictive. The extinction retrieval is better, though it underestimates the error in the PBL. Both fail to appreciate the error caused by the improper dead time correction.

Those figures also compare the retrieval to the methods outlined in Section 2.1.2. It is in mostly good agreement¹⁶ with them in the PBL and exhibits a lesser spread and error than the Ansmann solutions in the free troposphere (especially for the lidar ratio). The Fernald-Klett method gives equivalent answers when given the correct $B(R)$. Such comparisons will be considered more thoroughly in Chapter 5.

2.6.2 Parameter error

In real retrievals, there will be some error in the estimates of the model parameters \mathbf{b} . This additional uncertainty can be factored into the retrieval by extending the measurement covariance, \mathbf{S}_ϵ , to cover all sources of error extraneous to the retrieval,

$$\boldsymbol{\epsilon}_y = \boldsymbol{\epsilon} + \mathbf{K}_b(\mathbf{b} - \hat{\mathbf{b}}) + \Delta\mathbf{f}, \quad (2.6.1)$$

¹⁶For a fair comparison, (2.1.15) is averaged over 300 m to be equivalent to the effective resolution of the retrieval.




Figure 2.16: Covariance matrix for retrieval of $\ln \beta$ and B from case 1. **(a)** Standard deviation of $\ln \beta$, estimated by the diagonal of \hat{S} , with the *a priori* value shown dashed. **(b)** Standard deviation of B . **(c)** Autocorrelation matrix for $\ln \beta$ (\hat{S} normalised by its diagonal). **(d)** Intercorrelation matrix for $\ln \beta$ and B (x and y , respectively). **(e)** Autocorrelation matrix for B .

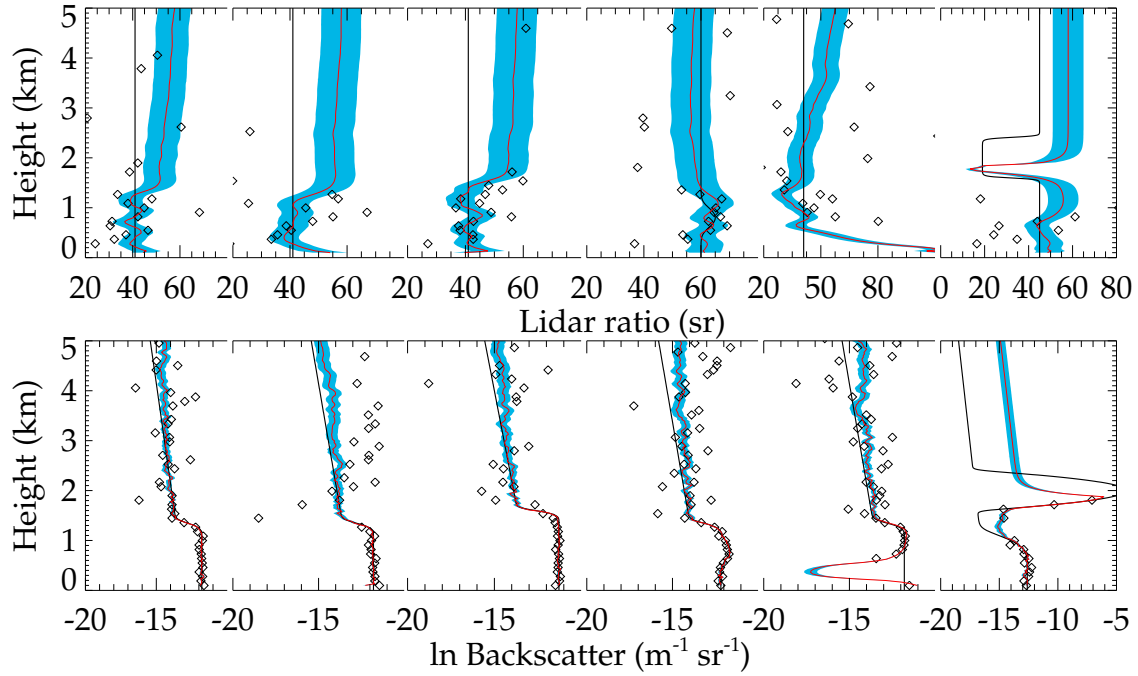


Figure 2.17: Result of the lidar ratio configuration showing its error compared to that of the Ansmann method (diamonds). True profile shown in black.

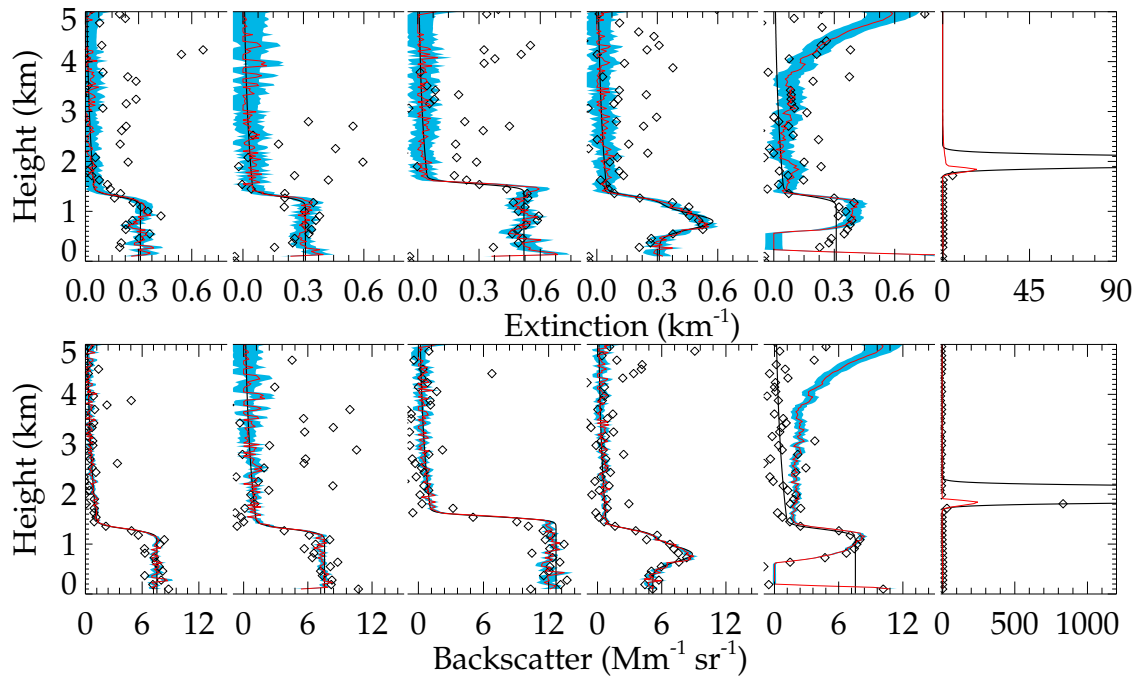


Figure 2.18: As fig. 2.17, but for the extinction configuration.

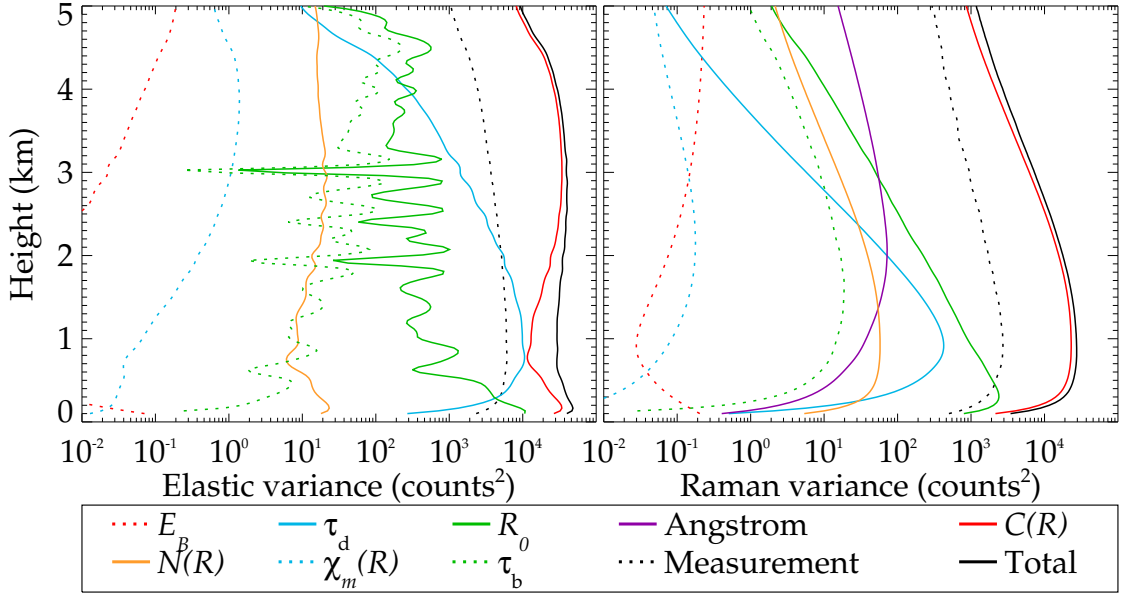


Figure 2.19: The contributions of model parameters to the total retrieved variance.

where $\mathbf{K}_b = \partial \mathbf{F} / \partial \mathbf{b}$ (see Section A.2); $\hat{\mathbf{b}}$ is the best estimate of the true parameters \mathbf{b} ; and the last term describes any inability of the forward model to describe the true state.

Concentrating on only the parameter error for now, this can be implemented by replacing all occurrences of \mathbf{S}_ϵ with,

$$\mathbf{S}_y = \mathbf{S}_\epsilon + \mathbf{K}_b \mathbf{S}_b \mathbf{K}_b^T. \quad (2.6.2)$$

This significantly increases the computing cost of the retrieval, as \mathbf{S}_y must be inverted for each iteration, which was not necessary when it was a diagonal matrix. A reasonable approximation is to only evaluate \mathbf{S}_y in the first and last iterations, but the full inclusion will be considered here.

This requires an estimate of these errors. The Ångström coefficient can vary quite significantly, but is commonly accepted to lie in the range 0.6–1.4, such that an error of 0.4 is not unreasonable [Klett, 1985]. Radiosondes measure pressure and temperature with an accuracy of 0.5 hPa and 0.2 K, from which an error in number density and χ^m of 0.5 % is expected. The standard deviation of the data used to estimate E_B can be easily derived.

The remaining parameters are somewhat more complicated as they must be calculated from some calibration procedure. These will be discussed in the next two chapters, but for the purpose of demonstration figure 2.19 shows the impact each parameter on the total variance, assuming errors in C of 10 % and τ_d of 1 ns;

such errors dominate the total¹⁷. For the impact of uncertainty in C to be of a similar order to the measurement error, it needs to be known to better than 2% — an unrealistic expectation. However, applying more realistic error estimates does not appear to greatly affect the result of the retrieval, such that this term only increases the overall error.

The dead time is more troublesome for elastic measurements as it introduces significant correlation to the errors around the peak of the measurement. If the dead time error is greater than 0.1 ns, this significantly reduces the information content available there. This correction clearly requires significant attention to permit suitable retrievals when the nonlinearity is so great. A system with a lower τ_d will have a greater tolerance for its estimation.

The laser energy E_L behaves similarly to C but is considered separately as it can change significantly with time whilst the calibration function should be fairly consistent. In the data to be considered later, the laser energy was often not accurately measured and so must be deduced directly from the data. It will therefore be retrieved as part of the state vector.

2.6.3 Further errors

The settings of the retrieval that have no bearing on the forward model should not affect $\hat{\mathbf{x}}$. The initial value of Γ_i alters the number of iterations required to converge as it drives the size of each step in state space. A value of 100 appears to be optimal in most cases and $\hat{\mathbf{x}}$ appears to be independent of that choice provided it is not too large or small. Similarly, convergence thresholds of 10^{-4} on change in cost or step and 0.1 on error were selected as the highest order for which $\hat{\mathbf{x}}$ is not greatly affected. The minimum retrieved height does have a small effect on the retrieval in the first few bins, so 100 m was chosen as low enough that these effects are concentrated within a region where parameter errors will be large regardless (see fig. B.5).

Forward model error is defined in Section 3.2.3 of Rodgers [2000] as,

$$\mathbf{G}_y[\mathbf{f}(\mathbf{x}, \mathbf{b}, \mathbf{b}') - \mathbf{F}(\mathbf{x}, \mathbf{b})], \quad (2.6.3)$$

where $\mathbf{G}_y = \partial \hat{\mathbf{x}} / \partial \mathbf{y}$, the sensitivity of the retrieved state to the measurement, and \mathbf{f} is the exact, true profile including any processes the forward model \mathbf{F} may not

¹⁷Interestingly, other than C , all errors for the Raman channel are an order of magnitude smaller than measurement error.

describe. This systematic error is generally difficult to estimate as, if f were known, it would most likely be used as the forward model instead.

There are some processes that are clearly not included within the current forward model. Multiple scattering has been neglected as it is mostly important for lidars with a wide footprint, such as space-based system, or for observations within clouds, where this algorithm is already known to perform poorly for other reasons. Though lidar-appropriate numerical models of multiple scattering exist [Eloranta, 1998], this is considered an area for future work if retrievals in clouds are desired.

There is also the small difference between the bin-averaged backscatter that is sampled and the true backscatter defined by Mie theory, for which the error can be evaluated with (2.6.3). It is, of course, greatest in the entrainment layer (or at any other sharp gradient) and even then is at most 1 % of the retrieved error for case 1.

The models of the calibration function and detector nonlinearity are idealised versions of the truth. A rough estimate of these contributions can be produced by considering alternative models (such as a paralyzable model of a Donovan correction). For case 5, these are over 100 times larger than the other errors, though that is still insufficient to make the profile consistent with the truth. This again demonstrates that great care will have to be taken over the dead time correction to ensure sensible retrievals and error estimation.

It should also be noted that none of the errors described would cause the error bars of figure 2.17 to coincide with the truth at the top of the profile. This is because that region has not deviated from the *a priori* state. Increasing the *a priori* variance, reassuringly, doesn't greatly affect the retrieved profiles. Hence, when applied to real data in Chapter 5, the *a priori* error will be increased to 40 sr to more clearly highlight where the retrieval has added information.

2.7 Aside

It is very simple to briefly inspect the potential of the retrieval to work with only one channel by artificially increasing the measurement error on a channel. The results, figure 2.20, are surprisingly encouraging. A lone lidar ratio retrieval, as might have been expected from the averaging kernels (fig. 2.15), hardly diverges from its *a priori*, but a lone extinction retrieval does a fairly respectable job.

The ability of either configuration to derive backscatter from one channel is dependent on the *a priori* lidar ratio, as the retrieval has insufficient information

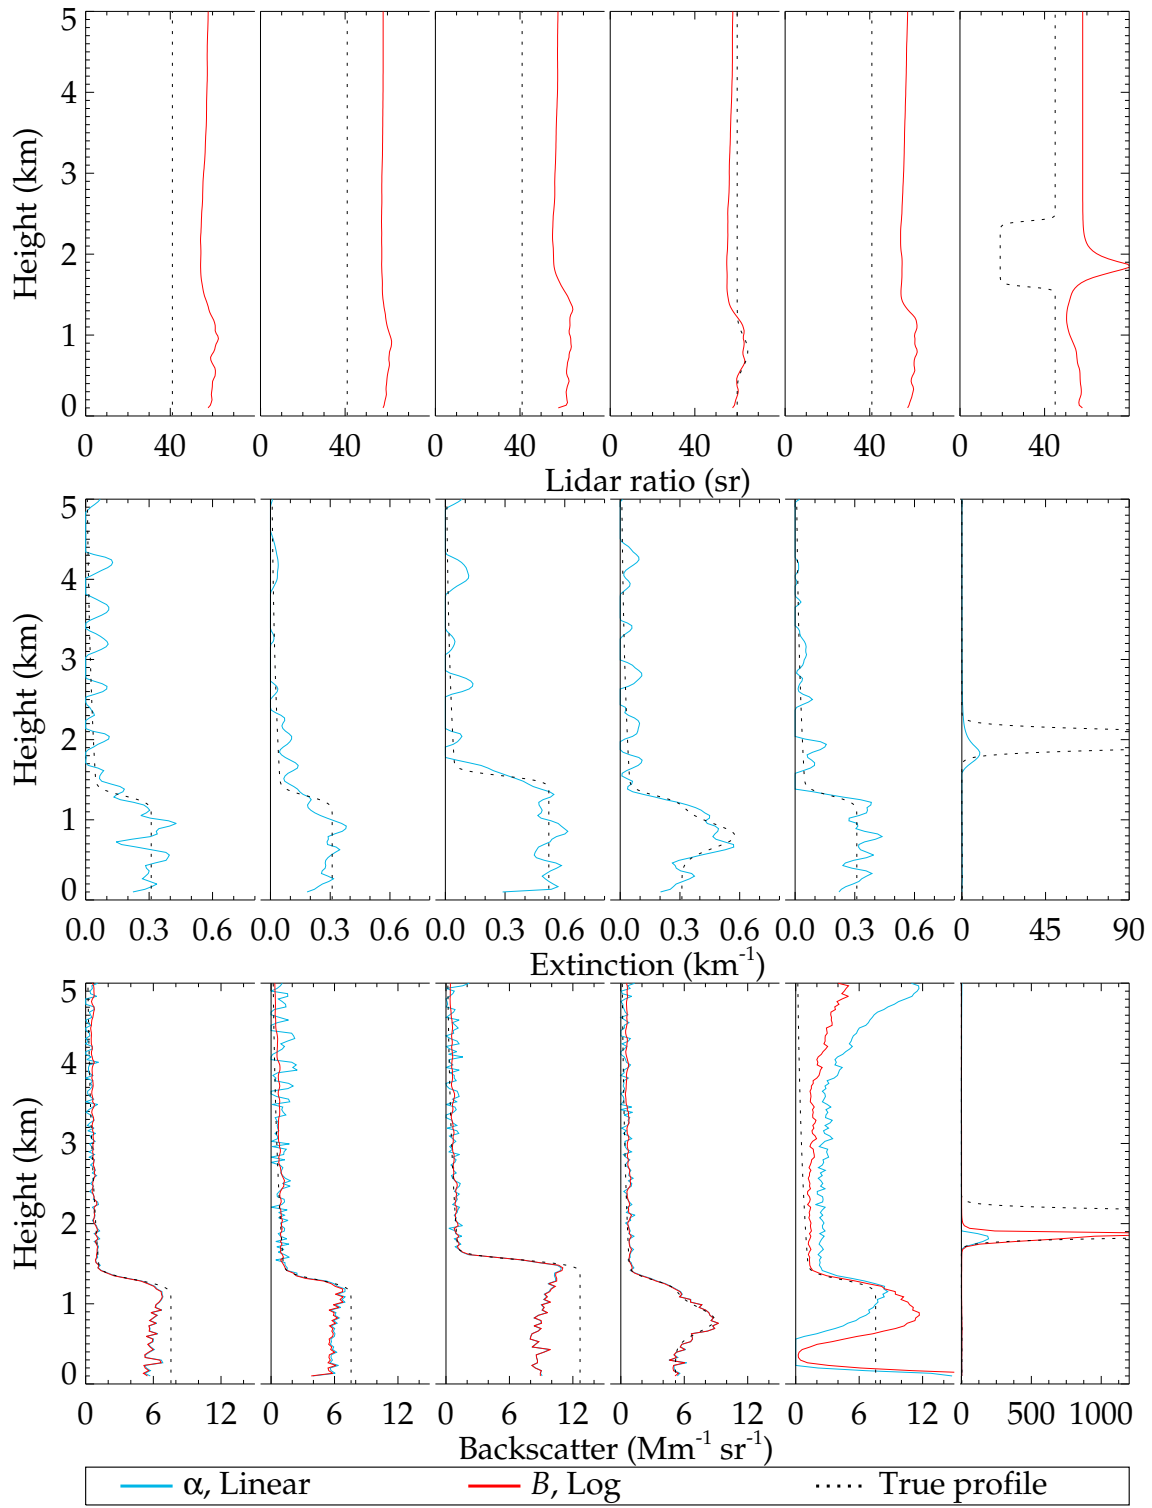


Figure 2.20: A demonstration of one-channel retrievals.

to alter that. Case 4 coincidentally has a fairly good initial guess and so does an impressive job fitting the aerosol layer. Clearly, there is some value to the algorithm beyond the particulars of a Raman lidar, though they are not further explored in detail here.

2.8 Summary

An optimal estimation retrieval scheme for aerosol scattering properties is proposed, using the lidar equations as a simple forward model. The *a priori* state and covariance matrix are based on the properties of aerosol outlined in the OPAC model. Scattering is assumed to be vertically correlated between bins, decaying exponentially over a scale height of 100 m. This is smaller than observed by balloon-borne measurements but ensures that the PBL and free troposphere do not greatly influence each other.

The details of the forward model were refined by considering their impact on the retrieval's ability to process simulated data. In particular, the state of the atmosphere can be described at each height by the aerosol backscatter and either the extinction or lidar ratio. If the lidar ratio is to be retrieved, the logarithm of the backscatter should be retrieved. This accurately reproduces common features in the PBL and gives a smooth lidar ratio profile with an effective resolution of about 300 m for the data considered. It loses sensitivity rapidly as the SNR decreases, producing biases in the free troposphere that are not adequately represented within the retrieved errors. If extinction is retrieved instead, it and backscatter should be retrieved linearly with a correlation assumed between them (95 % here, though more investigation of this value is necessary). This produces noisier profiles than the other configuration, especially in its derivation of the lidar ratio, but it does not lose sensitivity as aerosol abundance decreases in the free troposphere.

Hence, both configurations offer certain advantages and disadvantages. The smoothness and lower error of the lidar ratio configuration in the PBL are appealing, but due to a greater reliance on the *a priori* state, it is likely inferior to the extinction configuration. Both will be used to evaluate real data in Chapter 5 to further explore their concurrence with existing methods of analysis, especially as significant differences between their results would indicate a failure of the forward model. Regardless, the quality of any retrieval is highly dependent on the quality of the data given to it. Four parameters in particular — the dead time correction, the calibration function, the laser energy, and the vertical offset — need to present

minimal error for the retrieval scheme to consider the input data to have any useful information content.

Chapter 3

Assessment of model parameters

The previous chapter considered only simulated data and did not address how the various parameters of the forward model would be determined in practice; this chapter will outline such techniques. In particular, it will consider the degree to which these terms can be evaluated for the RACHEL and CUV systems some time after data were collected and additional calibration measurements cannot be performed. The calibration function will be considered separately in the next chapter.

3.1 Environmental terms

3.1.1 Pressure and temperature

An estimate of the number density profile $N(z)$ is required in order to determine the degree of molecular and Raman scattering. This is derived from coincident radiosonde launches or a standard atmosphere constrained by surface pressure and temperature [Weitkamp, 2005, chapter 4]. The discussions of Section 2.6.2 showed that the error due to these terms is fairly unimportant, such that the standard atmosphere is likely sufficient for most purposes.

The most common standard atmosphere [NOAA [1976]] assumes a constant lapse rate of $\Gamma = -6.5 \text{ K km}^{-1}$ through the troposphere such that,

$$T(z) = T(0) + z\Gamma \quad p(z) = p(0) \left[\frac{T(z)}{T(0)} \right]^{-\frac{gm_A}{k_B\Gamma}} \quad N(z) = \frac{p(z)}{k_B T(z)}, \quad (3.1.1)$$

for $z < 11 \text{ km}$ where m_A is the mean molecular mass of dry air and g is gravitational acceleration. The stratosphere (11–20 km) is assumed isothermal such that pressure decreases exponentially with a scale height of $k_B T(11 \text{ km})/gm_A$. The sur-

face pressure and temperature are estimated by the mean observed with electronic sensors¹ during the lidar's averaging period.

Where radiosonde measurements are available, though, it is reasonable to use them. Within the UK, the Met Office launches two sondes a day from ten sites across the country (fig. B.6). Limited resolution forms of these profiles are released for academic research [UK Met Office [2011]], such that any area of England is generally within 200 km and 24 h of a launch. Kitchen [1989] investigated their ability to represent the atmosphere at distant locations by evaluating the accuracy with which a profile could approximate those adjacent to it in space and time. He found that, in the absence of a frontal system, the resulting errors were approximately 2 % for sondes launched within 200 km and 24 hours. However, this is not applicable to the PBL, where small scale variations perturb the temperature profile. A detailed analysis of these errors has not been found and the desire to minimise them inspires some studies to make dedicated radiosonde launches as often as every 30 minutes [e.g. Ferrare et al., 1998; Schmid et al., 2006]. Indeed, many studies using Raman lidar consider a ratio of two Raman channels to eliminate, amongst other things, the impact of N [Ansmann et al., 1992].

It is necessary to interpolate radiosonde profiles onto the lidar's range grid from their native pressure coordinate. Measurements of altitude are linearly interpolated against $\ln p$ to give the height of each pressure level. The pressure at each lidar range bin is then estimated by the reverse interpolation of $\ln p$ against height. Temperature is interpolated linearly against height. Surface observations, where available, are also included as a data point in the interpolation (the impact of which has not been investigated). When two launches occur within 24 hours, the profile is assumed to evolve linearly over time. Otherwise, the profile is assumed to vary linearly towards a standard atmosphere.

Figure 3.1 illustrates the interpolation. Errors (blue) are estimated assuming accuracies of 0.5 hPa, 2 K, and 60 m, which are dominated by representativeness error. The pressure profile does not vary significantly. Temperatures are consistent through the free troposphere but, as expected, show greater variation in the PBL. Also, the linear interpolation performs poorly in the stratosphere. However, this thesis will consider only the first few kilometres of the atmosphere and so shall ignore this inaccuracy.

¹Further details are provided at http://badc.nerc.ac.uk/data/chilbolton/met_sensors.html for CFARR and <http://www.atm.ox.ac.uk/group/eodg/weather/station.html> for Oxford.




Figure 3.1: Illustration of variability in radiosonde profiles from Larkhill (51.2° , -1.8°) at 11:46 on 21 Apr 2010. Pressure is plotted as a dashed line and appears on the left of each plot. **Left:** Raw data plotted as crosses over the interpolated profile (error in blue) compared to a standard atmosphere in red. **Centre:** All profiles from Larkhill within 24 hours of the previous plot. **Right:** Profiles from Herstmanceux, Shoeburyness, and Watnall launched at a similar time.

3.1.2 Ångstrom coefficient

Following Ansmann et al. [1990], the Ångstrom coefficient is commonly assumed to equal unity for aerosol particles, as in (2.1.14). The previous chapter confirmed that the impact of this assumption is small but becomes increasingly significant with height (fig. 2.19). Its value can be estimated when aerosol optical thickness is measured at several wavelengths. The AERONET station at CFARR [Woodhouse and Agnew, 2011] routinely measures χ_∞ at 340, 380, and 440 nm, from which the Ångstrom coefficient may be estimated by,

$$\kappa = \frac{\ln [\chi_\infty^{(0)} / \chi_\infty^{(1)}]}{\ln [\lambda^{(1)} / \lambda^{(0)}]}. \quad (3.1.2)$$




Figure 3.2: Observed distribution of the 340–440 nm Ångström coefficient observed by the AERONET station at CFARR during 2006–10.

Figure 3.2 presents a histogram of these measurements over recent years, showing that the previous chapter’s estimate of 1.0(4) accounts for 80 % of the observations. Hence, a value of unity will continue to be assumed for all retrievals.

3.2 Instrumental effects

3.2.1 Vertical offset

Data acquisition in a lidar will generally be triggered to begin collection by a pulse from the laser’s Q-switch. That pulse does not necessarily immediately increase the laser cavity’s Q-factor to begin lasing and there will also be some distance for the pulse to travel through the optical system, such that the first bin of the output will not correspond to integration from $R = 0$ to $\frac{1}{2}c\tau_b$. It will also average over some distance before the laser pulse. Hence, the range axis should be expressed in terms of a vertical offset R_0 ,

$$R_i = \frac{ic\tau_b}{2} + R_0. \quad (3.2.1)$$

The offset is ideally determined through laboratory testing — generating a trigger pulse, followed at a known interval by a signal pulse, and monitoring the bin in which the detector registers that signal. Its value shouldn’t change significantly such that a single estimation should suffice for some time.

If an estimate of the offset is desired without dismantling the optical system, it is possible to fit the profile from one lidar to a simultaneous observation by a

properly calibrated lidar. This is most successful when there is a clear feature in both, such as a thin cloud. Knowing the bin width, each profile is interpolated onto a fine range axis and the vertical offset that maximises the correlation between the two is sought. The correlation should be evaluated only for those bins in the region around the feature as otherwise the fit will be dominated by the peak of CR^{-2} . The accuracy of such a fit is limited to tens of metres by the size of the range bins, the quality of the reference lidar's calibration, and a myriad of other factors.

3.2.2 Laser energy

In most lidar systems, a small fraction of the laser beam is directed onto a power meter to monitor the laser energy, E_L . This was regrettably neglected from the RACHEL system and so, as mentioned in the previous chapter, it is intended to retrieve the laser energy. However, it is useful to be able to estimate its value, if only to act as a first guess for the retrieval. Many options have been explored and, though none are completely satisfactory, the most useful product considers the first few bins of signal in the Raman profile. There, the return can be approximated as,

$$E^{\text{ra}}(R) \simeq \frac{E_L C^{\text{ra}}(R, \epsilon)}{R^2} \frac{p_0}{k_B T_0} + E_B, \quad (3.2.2)$$

where p_0 and T_0 are the surface pressure and temperature. The average of this expression for the first three bins of CUV profiles is plotted against measurements by its power meter in figure 3.3. Though there are clearly issues with this estimation, there is sufficient correlation for it to serve as a first guess for the laser energy of the RACHEL system.

3.2.3 Background counts

There are many sources of dark current in a PMT that contribute to the background count rate. This work neglects the differences between them, simply estimating a total background count rate from the top of each profile. For RACHEL, this should be accurate as the profiles view beyond 30 km, where the laser-induced signal is certainly negligible². The CUV is more problematic as it reports to only 13 km, where a small, but significant, laser-induced signal remains. A few bins are observed before the laser pulse and the mean of these correlates almost exactly with the mean from the last 100 bins. Hence, the latter shall be used due to its

²In particular, the top 500 bins are used.

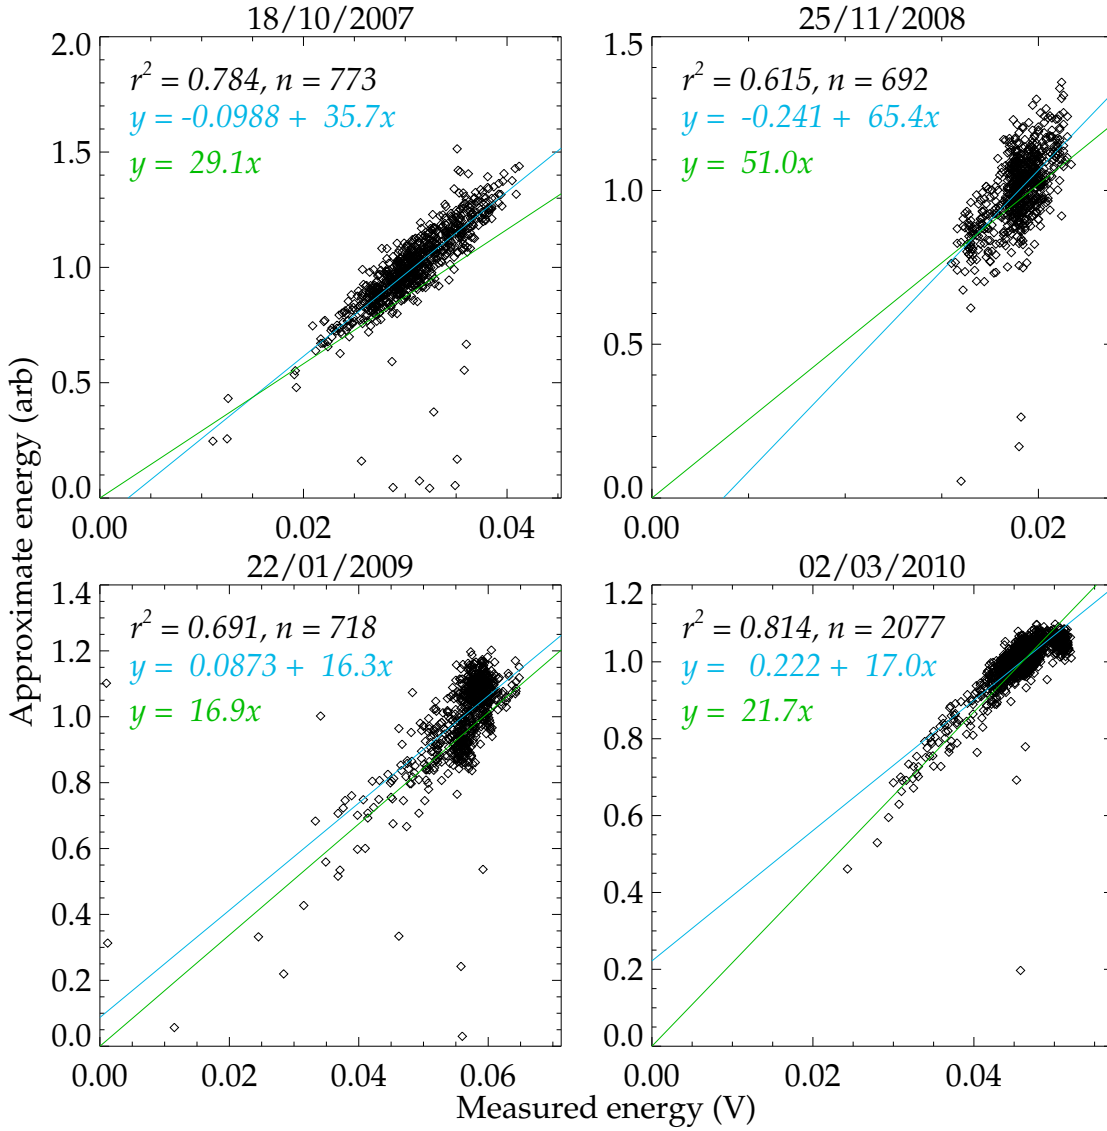


Figure 3.3: Performance of the laser energy approximation (3.2.2) on four typical days of observation with the CUV, showing the estimate (normalised by its average) against the measurement of a power meter. **Blue:** Linear fit to the data. **Green:** A linear fit constrained to intersect the origin.

lower random error, though the possibility of a systematic overestimation of E_B should be remembered.

3.2.4 Detector nonlinearity

Section 1.3.3 briefly outlined models for the nonlinearity of a photon-counting PMT. In the laboratory, many methods exist to measure this directly, such as observing the signal from two sources separately and then together [Evans, 1955;

Müller, 1973]. For a lidar, though, it is more practical to estimate the dead time from backscattered profiles rather than a disassembled instrument. Whiteman et al. [1992] is a rare example of a published methodology, proposing that two profiles are collected in succession — one normally and one with a 10 % neutral density (ND) filter placed before the PMT. The nonparalyzable correction (1.3.23) is applied to both for a variety of dead times and a ratio of the two profiles formed. The value of τ_d is estimated by that which “yields curves that are most nearly constant with altitude,” which is presumably assessed visually. This method gives τ_d to ± 1 ns [Whiteman, 2003]. Goldsmith et al. [1998] applies a similar technique but specifies that it uses “a nonlinear least-squares fitting routine to match the two profiles with two free parameters, the density of the ND filter, and the resolving time.”

By collecting a third profile — attenuated by an additional, identical ND filter — an analytic calculation of the dead time can be performed,

$$\varphi^{(0)} = \frac{E}{1 + \tau_d E} \quad \varphi^{(1)} = \frac{\mathcal{F}E}{1 + \tau_d \mathcal{F}E} \quad \varphi^{(2)} = \frac{\mathcal{F}^2 E}{1 + \tau_d \mathcal{F}^2 E} \quad (3.2.3)$$

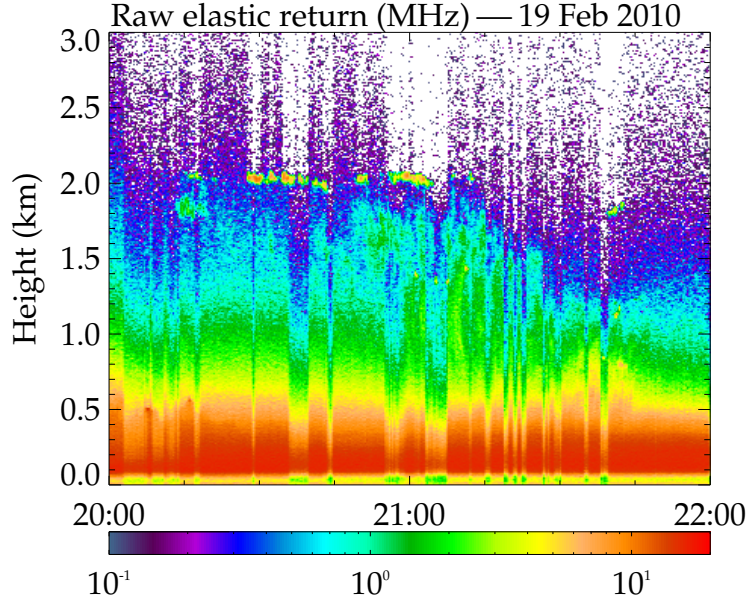
$$\therefore \tau_d = \frac{\varphi^{(1)} - (\varphi^{(0)}\varphi^{(2)}/\varphi^{(1)})}{\varphi^{(0)}\varphi^{(1)} - 2\varphi^{(0)}\varphi^{(2)} + \varphi^{(1)}\varphi^{(2)}}, \quad (3.2.4)$$

where \mathcal{F} is the transmission of the filter.

Alternatively, an increasing number of lidar systems now log both the analogue and photon-counting profiles from the detector. In most studies, these profiles are ‘glued’ together to form a single measurement using the analogue data near the surface where count rates are high and the photon-counting elsewhere for its greater SNR. A particularly noteworthy example of this is Newsom et al. [2009], which develops a regression scheme to determine the dead time, analogue offset, and proportionality constant required for this ‘gluing.’ It observed a diurnal variation of these values, confirming the importance of regular monitoring of the calibration of a lidar.

These methods, though, were of little use during the latter half of 2010. After its deployment at CFARR, the RACHEL system was non-operational for a year, making it impossible to perform the measurements required for such techniques. Thus, a new algorithm was required which could be applied to existing data. As already mentioned, during that campaign there were significant problems with laser stability. Figure 3.4 illustrates these frequent, sharp changes in the laser energy. This is normally problematic, but for the purposes of estimating the dead time provides exactly what is required — successive profiles at different energies.

Figure 3.4: An example of laser energy variations in RACHEL showing the raw return in MHz of the elastic channel observed from 20:00–22:00 on 19 Feb 2010 over 10 s integrations. White indicates a measurement of zero.



The span is not as great as the 10 % used by Whiteman et al. [1992], but had to suffice until the system was again operational. Hence, it was proposed to develop a retrieval of the nonlinear correction for a PMT from two successive lidar profiles, assumed to observe identical atmospheres.

3.2.4.1 Forward model

This problem requires a fairly simple forward model. The state vector consists of the true laser backscattered profile $E(R)$; the energy of the second profile relative to the first \mathcal{F} ; and the parameters of the dead time model. The true background count rate E_B can be deduced from that observed by applying the nonlinear correction within the forward model. It will thus be considered a parameter.

Two forms of the correction are considered, where the superscripts H and L denote the high and low energy profiles:

- The nonparalyzable correction where,

$$\varphi_i^H = \frac{M\tau_b(E_i + E_B)}{1 + \tau_d(E_i + E_B)} \quad \varphi_i^L = \frac{M\tau_b(\mathcal{F}E_i + E_B)}{1 + \tau_d(\mathcal{F}E_i + E_B)}. \quad (3.2.5)$$

- The correction of Donovan et al. [1993] where,

$$\begin{aligned} \varphi_i^H &= M\tau_b(E_i + E_B) \exp[-\tau_d(E_i + E_B)] \\ &\quad \times [P_1 + P_2\tau_d(E_i + E_B) + P_3\tau_d^2(E_i + E_B)^2] \end{aligned} \quad (3.2.6)$$

$$\begin{aligned} \varphi_i^L &= M\tau_b(\mathcal{F}E_i + E_B) \exp[-\tau_d(\mathcal{F}E_i + E_B)] \\ &\quad \times [P_1 + P_2\tau_d(\mathcal{F}E_i + E_B) + P_3\tau_d^2(\mathcal{F}E_i + E_B)^2]. \end{aligned} \quad (3.2.7)$$

The form of \mathbf{K} is summarised in Section A.5.

This can be extended to consider more than two profiles by simply adding an additional \mathcal{F} variable for each profile and repeating the form of φ_i^L .

3.2.4.2 Expected properties and distribution

The choice of an appropriate *a priori* presents an interesting question as there is not actually much known before the measurement is made. Lidar profiles tend toward a similar shape, but it is impossible to describe them more accurately without an assessment of the aerosol loading. Similarly, \mathcal{F} is not known here. A manufacturer's estimate is available for a nonparalyzable dead time correction, but is given to only one significant figure without any comment on its variations over time. As such, optimal estimation may not be the best retrieval scheme to apply to this problem. However, provided the *a priori* constraints are sufficiently weak, the technique should merely be a computationally inefficient form of a least-squares fit. Further, computer codes were readily available for an optimal estimation retrieval whilst they were not for a least-squares fit with errors in both the x and y variables.

Some knowledge of the lidar profile can be deduced from readily available data. Over 20,000 ABC profiles collected by the CUV on 28 days since 2006 were available for this research, from which 1,500 profiles were randomly selected to generate a covariance matrix, shown in figure 3.5. The analogue channel was considered because of its perceived linearity³.

The elastic covariance matrix is similar to the form of the backscatter *a priori* (with correlations dropping rapidly with distance) as the ABC is dominated by β . A fit of (2.3.1) to that matrix, shown on the right of fig. 3.5, gives a correlation scale height of 240 m. The Raman channel appears to present a very different form with significant correlations to some point around 2.5 km. The cause of this is clear from the variance (plotted left) — the Raman channel simply has too little signal available at most heights, the average SNR being less than 10. A more obviously diagonal matrix can be obtained by averaging the profiles over greater times but is not shown as it is desired to consider the data in their rawest form. For current purposes, (2.3.1) was fit only to data $R < 3$ km, giving a scale height of 3.2 km.

In the absence of better information, \mathcal{F} is assumed to take a value of 1 ± 1 . The dead time is taken to be 40(20) ns. For the Donovan correction, $P_1 = 1$ and

³The photon-counting channel gives similar results but with greater correlations in the elastic channel through the lowest 2 km due to its nonlinearity. Interestingly, clouds at the top of the PBL reduce correlations there.




Figure 3.5: Covariance matrices from 1,500 randomly selected CUV analogue profiles averaged over 20 s. **Top:** Elastic channel. **Bottom:** Nitrogen Raman channel. **Left:** Observed variance as a function of height. **Centre:** Observed correlation matrix. **Right:** Fit of (2.3.1) to observations.

$P_2 = P_3 = 0$ is assumed, corresponding to a purely paralyzable detector, with an uncertainty of unity on each term. The first guess of $E(R)$ is simply the highest energy profile with its background subtracted, without further correction. The laser energy is estimated using (3.2.2).

Six pairs of profiles, identical to case 1 of Chapter 2, were simulated with $\mathcal{F} = 0.1$ and various nonlinear corrections:

1. A nonparalyzable correction with $\tau_d = 29$ ns.

2. A nonparalyzable correction with $\tau_d = 50$ ns.
3. A paralyzable correction with $\tau_d = 8$ ns.
4. A paralyzable correction with $\tau_d = 30$ ns, though this has been adjusted such that the observed count rate does not decrease after reaching a maximum.
5. A Donovan correction with $\tau_d = 15$ ns, $P_1 = 0.9$, $P_2 = 0.08$, and $P_3 = 0.01$.
6. A Donovan correction with $\tau_d = 23$ ns, $P_1 = 0.6$, $P_2 = 0.3$, and $P_3 = 0.05$.

Profiles were simulated for both the elastic and Raman channels by assuming a constant lidar ratio of 41 sr.

3.2.4.3 Retrieval assessment

Figure 3.6 presents the retrieved nonlinear correction via the observed versus true count rate. Retrieval from both the elastic and Raman channel are considered (blue and green lines, respectively). The nonparalyzable configuration (top) converges close to the correct answer if the data were similarly afflicted, retrieving $\tau_d = 29.3(2)$ and $50.3(3)$ ns for cases 1 and 2. It is also close to (if not entirely consistent with) the truth in cases 3 and 4 until the detector saturates.

In the remaining cases, the retrieved correction is useful only at low count rates. This is more clearly shown by applying the retrieved correction to the observations and plotting them against each other. Perfectly corrected data would plot a one-to-one line, which cases 4–6 clearly fail to do for rates above 20 MHz. This is consistent with the wide-spread application of the nonparalyzable correction to lidar data, which is understood to apply only up to some maximum count rate, generally where the correction accounts for 10 % of the signal [Hakamata, 2006, chapter 6]. Hence, one could perform a series of retrievals, removing excessively bright measurements with each iteration, until a suitable fit is obtained or, preferably, design a forward model that somehow accounts for this upper limit.

The retrieval with the Donovan correction fails utterly, producing noticeably biased corrections (particularly case 3). Like the other configuration, it converges in generally less than 10 iterations to a solution with an anomalously low cost⁴. Unlike the other configuration, though, its solutions are highly sensitive to the first guess and other parameters; there simply is not enough information available to

⁴ ~ 0.5 as, regardless of the first guess, it fits E to φ^H almost exactly because that has the largest absolute error. An x - y fit with errors on each would probably be preferable.

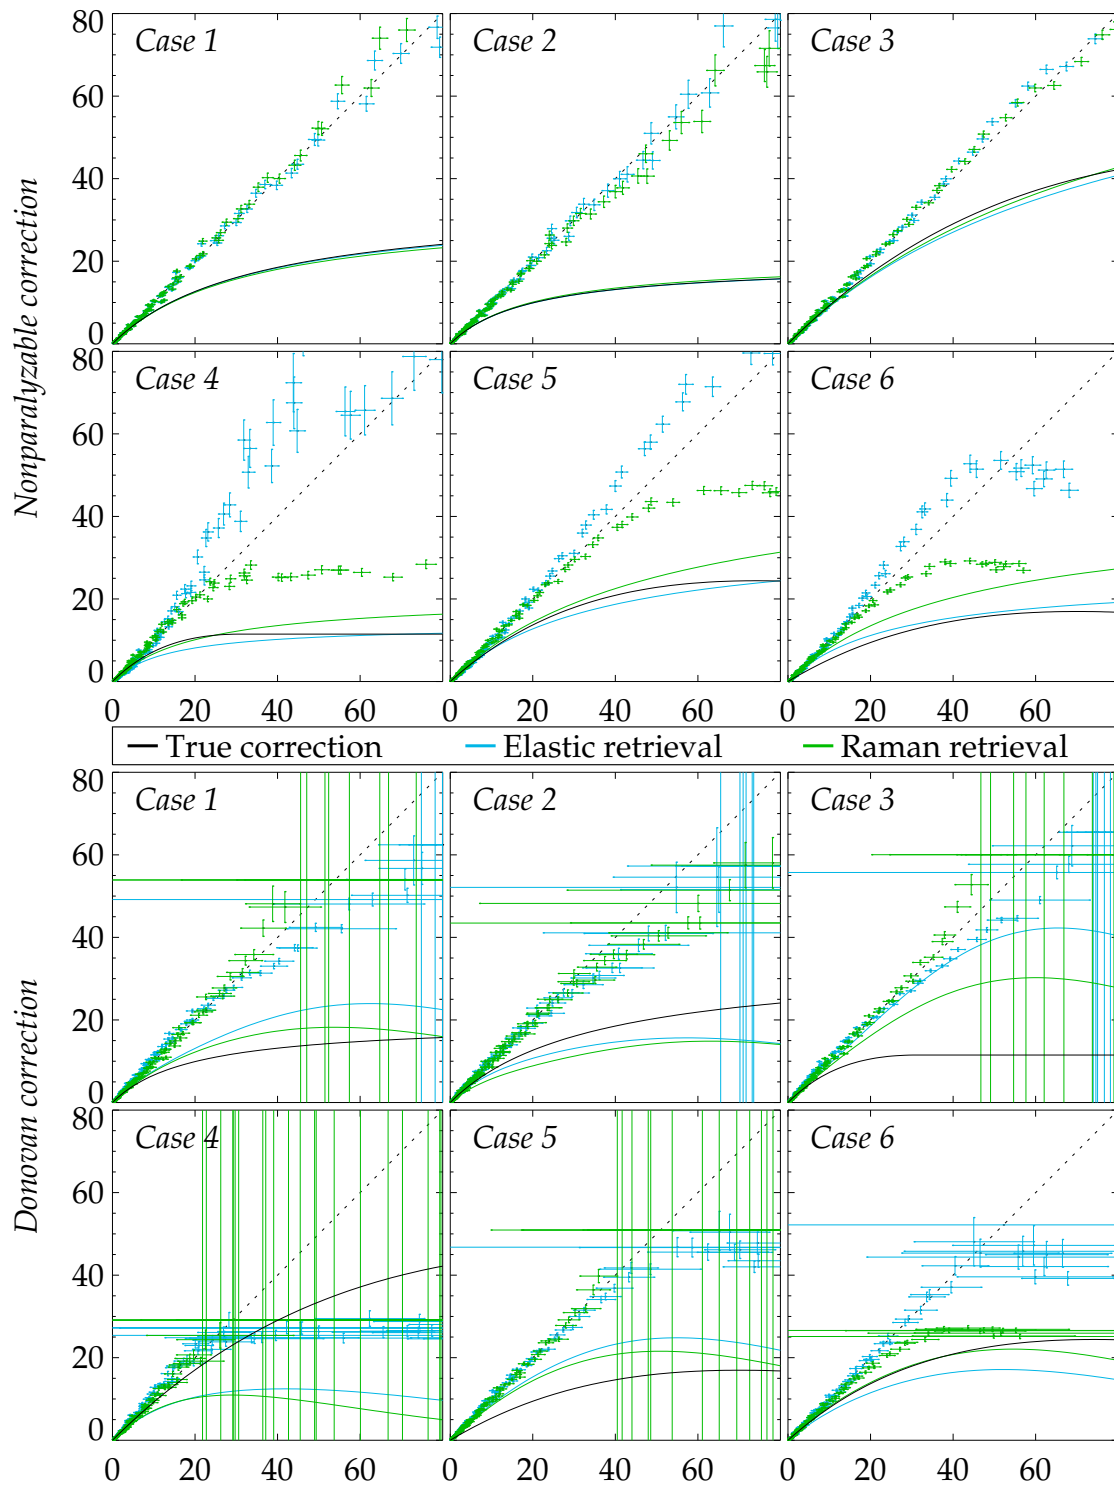


Figure 3.6: Performance of the dead time retrieval. **Top:** Retrievals using a non-paralyzable correction. **Bottom:** Retrieval using a Donovan correction. **Lines:** Observed vs. true count rate in MHz. **Crosses:** High vs. low energy profile after the application of the retrieved nonlinear correction. Bar length indicates propagated measurement error.

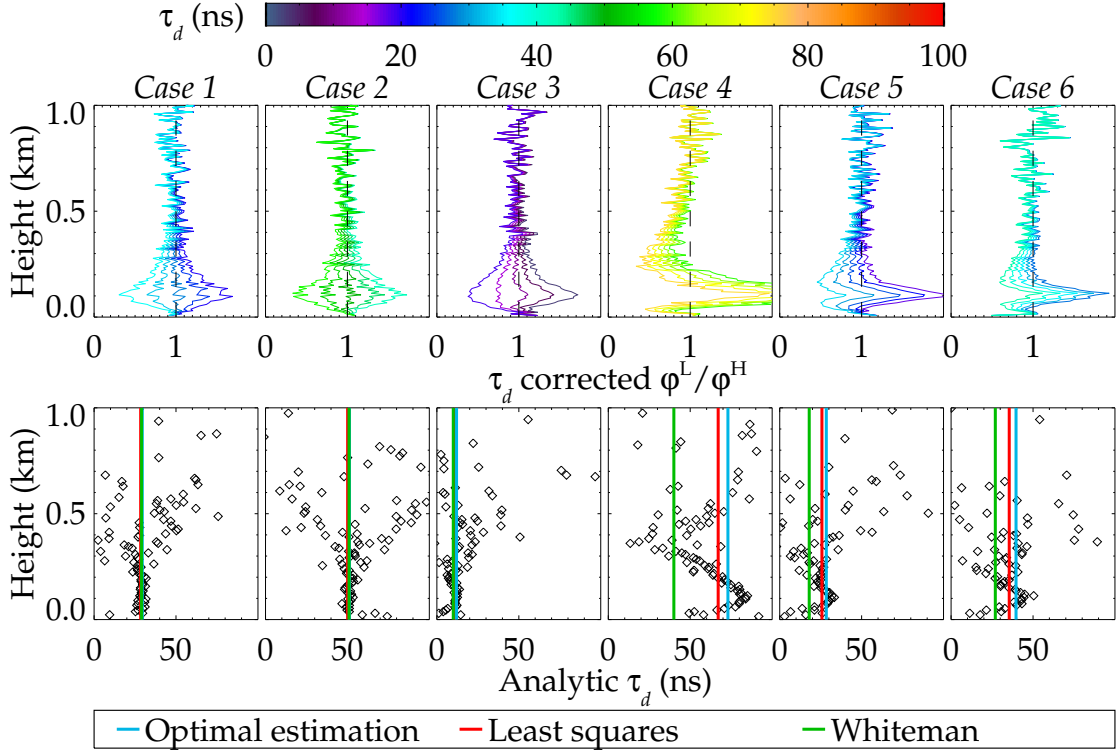


Figure 3.7: Comparison of the dead time retrieved for the elastic channel by various methods. **Top:** The ratio (3.2.8) for dead times near to that retrieved by a simple least squares fitting, colour indicating τ_d . **Bottom:** Evaluation of (3.2.4) for profiles with $\mathcal{F} = 0.1$ and 0.01 , also showing τ_d retrieved by the optimal estimation, least-squares fit, and Whiteman technique for comparison (blue, red, and green, respectively).

constrain the additional terms. At low count rates, where data are plentiful, there is little differentiating the various possible models and so these points contribute little information to the retrieval. For the Donovan correction, increasing the count rate also increases the retrieval error as the gradient of the correction curve becomes shallower, such that the most relevant data then suffer an excessive error. Hence, though the retrieval may coincidentally give a sensible answer, it does not give a consistent one.

Significant efforts were made to find enough information to perform a retrieval. For example, it could consider a series of profiles at different laser energies rather than just two. A highly attenuated profile over a longer integration time could be made to better constrain⁵ E . Alternatively, the laser energy could be measured.

⁵Following the previous footnote, if the integration time of the lowest energy profile is long enough that its absolute magnitude (in photon counts) is greater than that of the highest energy profile, it should dominate the retrieval of E . This would require an unusually stable atmosphere.




Figure 3.8: As fig. 3.6 for the elastic channel, but comparing the performance of the three methods outlined in this section.

Any practical version of these does not significantly alter the retrieved result as the original pair of profiles already sampled the region of low count rates sufficiently and these additions simply contribute more data there. Sampling higher count rates *in situ* is not possible as any sensibly designed lidar will ensure that the peak count rate rarely saturates the detector — both to minimise the impact of nonlinearity and to prevent damage to the detectors. Hence, it would likely *never* be possible to fit a Donovan style correction to lidar profiles in this manner without specifically designing the system for that purpose or disassembling the optical components to perform laboratory tests.

Thus, the best that can be done *after* a measurement campaign is to fit a non-paralyzable correction, assess the range over which this is valid, and reject excessively bright data (or at least consider it with great caution). Figure 3.7 compares the various methods of estimating the nonparalyzable dead time outlined at the start of this section. A least-squares fit⁶ of the high to the low energy profile was performed, retrieving \mathcal{F} and τ_d . This is compared to the method of Whiteman et al.

⁶Neglecting measurement error on φ^L .

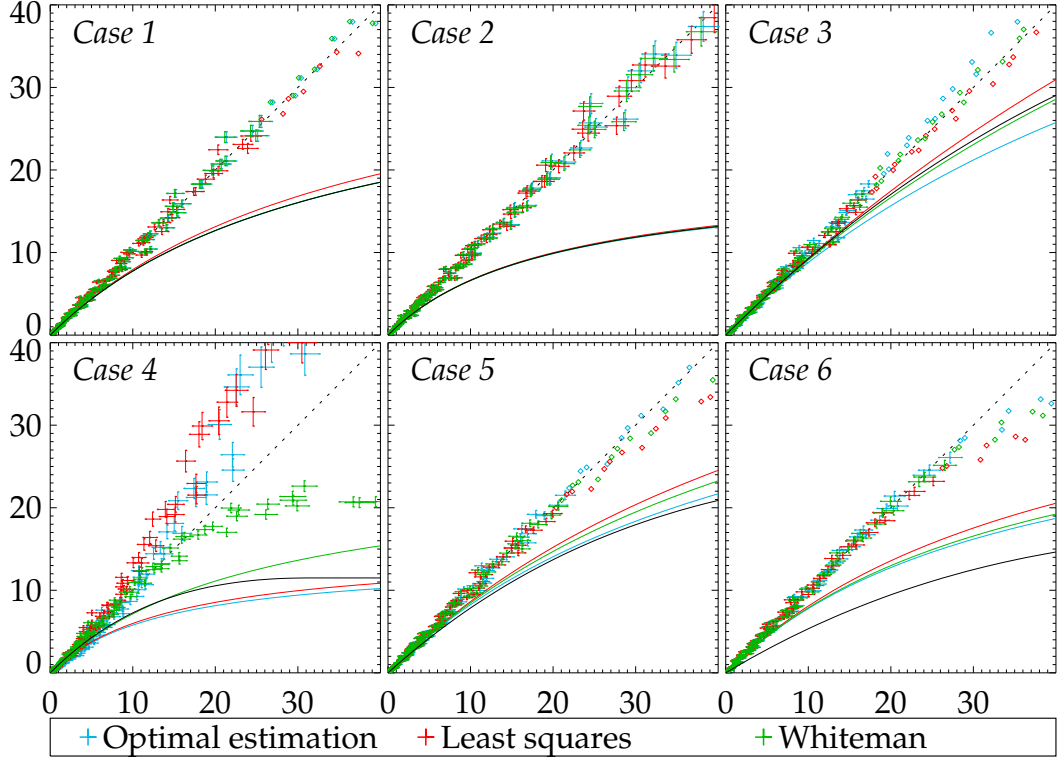


Figure 3.9: As fig. 3.8, but only fitting points < 15 MHz. Points not included in the fit are plotted as diamonds.

[1992] at the top of figure 3.7 by plotting the ratio,

$$\frac{\varphi^L(1 - \tau_d \varphi^H)}{\mathcal{F} \varphi^H(1 - \tau_d \varphi^L)}, \quad (3.2.8)$$

for the fitted values and those $\pm 4, 8$ ns for contrast. The two schemes agree in cases 1–3 (i.e. the centre-most profile is the most consistent with height), but they diverge in the more distorted cases, with the Whiteman scheme favouring smaller τ_d .

The bottom of figure 3.7 then considers the analytic formulation (3.2.4) from profiles with $\mathcal{F} = 0.1$ and 0.01 . The noise on these estimates is significant (mostly due to the 0.01 profile) and the results could be improved by increasing the integration time at the expense of increased error from atmospheric variations. Regardless, where the SNR is largest, the analytic estimates of τ_d are consistent with those retrieved by the other methods in cases 1–3. For cases 4–6, the optimal estimation retrieval returns consistently larger values than the least squares fitting, which itself is greater than Whiteman, while the analytic formulation presents no clear answer. These differences are highlighted in fig. 3.8. Though the more statistical fitting schemes are consistent with each other to fairly high count rates, they are




Figure 3.10: Application of the Whiteman dead time estimation to ten pairs of RACHEL profiles from 2010 where the laser power is observed to spontaneously change in the absence of obvious aerosol inhomogeneities.




Figure 3.11: Pulse linearity of the Hamamatsu R7400U PMT. Originally ETC [1998], reproduced with permission of Hamamatsu.

also biased high. The Whiteman correction gives an accurate result over a greater dynamic range.

This improved accuracy arises as, by selecting the τ_d correction “most nearly consistent” with height, the Whiteman scheme can clearly identify those regions where the correction is failing (100–200 m in fig. 3.7) and ignore those data. This could be applied to the fitting schemes by artificially inflating the error on data with a count rate in excess of some threshold. Figure 3.9 presents the impact of a 15 MHz threshold on the previous retrievals, which clearly improves the fits in all cases except 4. If the threshold is reduced to 10 MHz, that case also gives consistent results but that illustrates how an arbitrary cap is an unsatisfying solution as its application requires some *a priori* knowledge about the nonlinear correction. Hence, it appears that a truly independent retrieval is impossible — some judgement by a scientist is needed to constrain fitting to the range over which it is valid. Otherwise, there simply isn’t enough information available to make a consistent and accurate retrieval.

Hence, the Whiteman method is applied to RACHEL data. Ten pairs of profiles were selected from the CFARR deployment where the laser energy was perceived to suddenly change while no obvious features were observed for a period of five minutes before and after the change. The ratio (3.2.8) is shown in figure 3.10. The fits are broadly similar, giving $\tau_d^{\text{el}} = 32\text{--}40$ ns and $\tau_d^{\text{ra}} = 40\text{--}48$ ns (ignoring the last shown pair, where the profiles are too weak for a dead time correction to significantly alter the result). The nonparalyzable correction appears to be valid for all Raman observations, but saturation of the elastic channel between 100–300 m is evident.

3.2.4.4 *Ad hoc techniques*

Sadly, not even the Whiteman scheme can be applied to the CUV. Little information is available about the nonlinear correction for the CUV as CFARR does not apply one, routinely processing only the (assumed linear) analogue channel⁷. A technical summary of the PMTs presents a single plot, reproduced in fig. 3.11, of the ‘pulse linearity’ for the analogue channel. That is presented as a function of the current output, but as the CUV logs output voltage over an unknown resistance, this information is not overly useful. A single paper was found using the same detectors that applied a dead time of 4 ns [Kuang et al., 2011].

At the very least, it is possible to compare the two modes to assess the degree to which they may be linear. Figure 3.12 plots the analogue against PC returns for a random selection of data points from 2 Mar 2010⁸. Newsom et al. [2009] outlined that, for such measurements, the analogue voltage V should be a linear function of the true photon count rate φ_t ,

$$V(R) = c_0 + \frac{c_1 \varphi_m(R)}{1 - \tau_d \varphi_m(R)}, \quad (3.2.9)$$

where $c_{0,1}$ are constants and a nonparalyzable correction is applied. A linear fit of this to data less than 30 MHz where the impact of τ_d should be small is also shown. As the count rate increases, the PC measurement clearly diverges from the analogue, which should itself be linear as it encroaches on only 10 % of the dynamic range. It should be possible to derive a value of τ_d from these data to compensate for the lack of a direct measurement.

There are several possible ways to construct a nonlinear fit of (3.2.9) to these data. Applied to a single profile, regardless of the averaging period, there appears to be insufficient information to determine τ_d . The fit could process the entire data set at once, but as that involves tens of millions of points, it is deemed infeasible. Instead, the method of Newsom et al. [2009] is adapted: A value of τ_d is assumed. Then, for each profile⁹, a linear fit¹⁰ is made of the PC to analogue data with (3.2.9) using only data above background levels and the resulting χ^2 statistic is noted. The

⁷Data points where the analogue measurement exceeds 13 mV are rejected in the ABC product due to concerns of cross-talk between the channels and truncation by the ADC.

⁸This day was deemed representative of the entire data set as it captures day and night with clear and cloudy conditions.

⁹Clouds must be filtered from the data as the PC signal is beyond the range of a nonparalyzable correction. They are identified by their large positive gradient in the elastic channel.

¹⁰Using the error estimates of Section 5.1.2.




Figure 3.12: Nonlinearity observed in the CUV, plotting analogue against PC measurements for 2,000 randomly selected points from 2 Mar 2010. **Blue:** Raw data. **Green:** Data with a $\tau_d = 5.5$ or 4.5 ns correction for the Raman and elastic channels, respectively. **Black:** Fit of (3.2.9) to raw data < 30 MHz. **Red:** Same, but fitting the corrected data.




Figure 3.13: Average, normalised χ^2 as a function of τ_d from fitting (3.2.9) to four different days of data.




Figure 3.14: Variation of c_0 (top) and c_1 (bottom) from fitting (3.2.9) on 3 Mar 2010.

procedure is repeated for a range of τ_d and that giving the smallest χ^2 averaged over all profiles is taken as the best estimate of τ_d .

Figure 3.13 shows the χ^2 as a function of τ_d for four mostly clear days. The three from 2010 are in fair agreement, giving estimates of 5.5 and 4.5 ns to ± 0.5 for the Raman and elastic channels, respectively. The minima are shallower on 19 Apr 2010 as the large aerosol loading present ($\chi_\infty \simeq 1$) reduces the magnitude of Raman observations and so samples a smaller dynamic range. A day from several years earlier does not appear to require a nonlinear correction, partially because those data reach a maximum of only 35 MHz and partially because of difficulties with the analogue signal that will be discussed in the Chapter 5.

A moment should be taken to consider variations in the fit parameters $c_{0,1}$. Figure 3.14 shows the offset c_0 broadly follows a diurnal cycle whilst the gradient c_1 is more constant. This is consistent with the variations observed at night in figure 2 of Newsom et al. [2009], though that did not directly process daytime observations. The inclusion of errors in the fitting appears to be most important in the Raman

channel, where background noise is greatest, and produces smoother curves.

This algorithm is exactly the sort of *ad hoc* retrieval scheme that this thesis had proposed to remove from lidar analysis. Such techniques were necessary because of the numerous sources of variability in the data, illustrated by the impact of high aerosol optical thickness on the gradient of the χ^2 curve, and a lack of *a priori* information, extending even to the error of the analogue measurements¹¹. It would not be impossible to adapt the other methods to an optimal estimation scheme, automating the any judgements made, but that is a considerable task.

3.3 Summary

A number of methodologies were outlined in order to estimate the parameters needed for the forward model of Chapter 2 when applied to real data. Due to the nature of the available data, some consideration was given to the possibility of calibration using only existing observations. These were, for the most part, not satisfactory and so more traditional techniques requiring dedicated calibration measurements would be preferable, where possible.

The estimation of a correction for the nonlinearity of the PMTs was given detailed consideration. Despite its theoretical superiority, an optimal estimation retrieval scheme to fit the correction proposed by Donovan et al. [1993] to lidar profiles was unsuccessful as excessive model error eliminated the ability to discriminate between different state vectors. A number of techniques from the literature were compared and it was found that the least statistical, outlined in Whiteman et al. [1992], was also the most successful. Though the judgement of a scientist is clearly not an optimal approach, it provides an effective means to neglect data for which the simple nonparalyzable correction is not valid. The limited applicability of that correction, in addition to a general uncertainty in the community as to the exact form of a given PMT's response, likely explains why there is so little discussion of these issues in lidar literature. Many groups likely ignore the issue entirely by simply neglecting data that cannot be rectified with the usual correction. Despite significant effort, no success has been made in improving this situation. Laboratory measurements are likely the only accurate solution.

¹¹This could be estimated from the standard deviation of profiles around the measurement, but that estimate appears to be dominated by atmospheric variances in all data tested, providing a significant overestimate.

Using the Whiteman technique, the dead time of the elastic and nitrogen Raman channels of RACHEL were estimated to be 36(4) and 44(4) ns. This could not be applied to the CUV as no suitable measurements were available. An adaptation of the methods of Newsom et al. [2009] was used to fit PC to analogue measurements, finding dead times of 5.5(5) and 4.5(5) ns for the elastic and Raman channels.

Chapter 4

Estimation of the calibration function

The most significant source of error in the retrieval of Chapter 2 was the calibration function. It impacts the elastic profile in a similar manner to backscatter, such that it cannot be retrieved simultaneously without additional information. Further, it is a function of the precise alignment of the telescope and optical system, such that lab-based measurements are not necessarily representative of the instrument's behaviour in practice.

Ideally, the calibration function should vary smoothly and monotonically from zero to a constant value, which it maintains for all heights above that point. Assuming the instrument's operators can identify when it is so aligned, the difficulties of the calibration can be neglected by only considering the region of constant magnitude. Its value can be derived from observations of a clear sky (for which the extinction and backscatter are known) by averaging an algebraic inversion of the lidar equations (1.3.18) or (1.3.19) [Fernald et al., 1972; Reagan et al., 1989]. Long-term observations confirm that this value varies in magnitude, but over a period of days to weeks, which should be sufficiently slow that occasional calibration is sufficient [Agnew, J.L., personal correspondence].

However, when observations within the PBL are desired, it is necessary to estimate the calibration function. After briefly reviewing the extensive literature published on this topic, this chapter will detail and extend upon the work of Povey et al. [2012] to develop an optimal estimation retrieval for the calibration function of a Raman channel of a lidar.

4.1 Existing techniques

A wide range of techniques have been proposed to estimate the overlap function. When the arrangement of the system is well-known, the overlap function can be

determined directly using ray tracing techniques [Velotta et al., 1998; Berezhnyy, 2009] or approximated by an analytic formulation [Halldórsson and Langerholc, 1978; Ancellet et al., 1986; Kuze et al., 1998; Stelmaszczyk et al., 2005; Mao et al., 2012].

Otherwise, deriving the overlap function requires an estimate of the aerosol component of scattering. When the aerosol distribution is homogeneous, such as during light mist [Sasano et al., 1979] or observation of a well-mixed layer [Tomine et al., 1989], the signal from some point in the profile provides sufficient correction. Horizontal observations through a day-time mixed layer are commonly used for this reason [Berkoff et al., 2003]. It is also possible, with significant effort, to use an imaging system to compress the range of the system to facilitate its measurement in a laboratory [Hey et al., 2011]. In the absence of any additional data, it has been proposed that a polynomial fit to the range-corrected elastic profile in the region where complete overlap is expected can estimate the scattering lower in the atmosphere [Dho et al., 1997], though it is not clear why this should be successful.

For instruments in the infrared, the aerosol contribution can be measured with a coincident ceilometer as it does not suffer an overlap-based attenuation [Guerrero-Rascado et al., 2010]. More generally, wherever there are two instruments working at similar wavelengths, a well-calibrated system can be used to correct another by the ratio of their ABCs. If the capacity to scan the laser beam is available, it is possible to use multi-angle measurements to constrain the aerosol profile and derive a best-fit to the overlap function [Adam et al., 2007; Biavati et al., 2011].

For Raman lidar, by assuming the two channels share an overlap function, Wandinger and Ansmann [2002] showed that the calibration can be estimated from the difference between the backscatter profiles derived from each channel using the iterative scheme,

$$O^{\text{ra}}(R) = E^{(0)}(R)/E^{(\text{end})}(R) \quad (4.1.1)$$

$$E^{(i+1)}(R) = E^{(i)}(R) \left[1 + \frac{\beta^{\text{ra}}(R) - \beta^{\text{el},i}(R)}{\beta^{\text{ra}}(R) + \beta^{\text{m}}(R)} \right], \quad (4.1.2)$$

where β^{ra} is the backscatter profile obtained from application of (2.1.17); $\beta^{\text{el},i}$ is the backscatter profile obtained from the application of (2.1.10) to $E^{(i)}$; and 10–12 iterations are recommended.

The Wandinger scheme is fairly sensitive to the calibration applied in the calculation of β^{ra} . Figure 4.1 considers the solutions obtained from various randomizations of the same simulated profile. Their spread decreases as the SNR at the reference height increases, implying a minimum SNR of order 50 to obtain a sensible




Figure 4.1: Application of the Wandinger calibration method to data simulated for the RACHEL system with $\chi_\infty = 0.05$. Each panel shows analysis of 10 equivalent simulations over 10^5 , 10^6 , 10^7 , and 10^8 laser shots (left to right) and states the average SNR of the Raman channel at the reference height (10–11 km).

solution from this technique. The results also diverge from the truth with increasing χ_∞ as the scheme does not account for the residual transmission $\exp(\chi^{\text{ra}} - \chi^{\text{el}})$ in β^{ra} . Hence, in order to obtain a result for the low SNR cases, it is necessary to limit the bracket of (4.1.2) to take values $[0.1, 2.0]$. For low-powered lidar systems, the integration time required to achieve sensible results is often prohibitively long. Other schemes mentioned either require a dedicated measurement, very specific conditions, or both and so are not suitable to produce an estimate of the calibration for existing data. As such, a retrieval for the calibration function is proposed.

4.2 Forward model

By substituting (1.3.8) into (1.3.15), the calibration function can be written as,

$$C(\lambda, R, \varepsilon) = \frac{\eta(\lambda)A_0}{A_L(R)} \int \psi(R, \mathbf{r}) \xi(\lambda, R, \mathbf{r}) dA(R, \mathbf{r}) \frac{2}{c\tau_L} \ln \left[\frac{R \left(R + \frac{c}{2}\tau_b - \frac{c}{2}\tau_L \right)}{\left(R + \frac{c}{2}\tau_b \right) \left(R - \frac{c}{2}\tau_L \right)} \right]. \quad (4.2.1)$$

Taken with the lidar equations (2.2.10)–(2.2.12), there are four unknowns for each height bin and two measurements, which is an underconstrained problem. This

issue is avoided in most existing methods by considering circumstances where the scattering profiles are known.

If the Raman channel is considered alone, there are only two unknowns for each measurement — the calibration function and the extinction. To reduce the number of unknowns further requires an analytic formulation of the profile in terms of a few parameters. By limiting observations to circumstances in which such formulations are a reasonable approximation of the truth, a forward model for the calibration function can be developed.

4.2.1 Extinction

When the PBL is well-mixed, aerosol extinction generally tends towards a common profile — constant above and below the entrainment layer and sharply decreasing within it [Wiegner et al., 2006; He et al., 2008; Maurya et al., 2010]. Steyn et al. [1999] fit an error function,

$$\alpha(z) = \alpha_+ + \alpha_- \operatorname{erf} z_\sigma, \quad (4.2.2)$$

to that profile as a means to identify PBL height z_{BL} and entrainment layer thickness σ_{BL} , where z is vertical height; $z_\sigma = \sigma_{BL}^{-1}(z_{BL} - z)$; and $\alpha_\pm = \frac{1}{2}(\alpha_{BL} \pm \alpha_{FT})$ for a mean extinction of $\alpha_{BL,FT}$ within the PBL or free troposphere, respectively. Though useful in investigating PBL height, this is clearly not a suitable model of extinction as it integrates to infinite optical thickness. That can be addressed by introducing an exponential decay above the PBL, as expected in the OPAC model and used in (2.3.2),

$$\alpha(z) = \begin{cases} \alpha_+ + \alpha_- \operatorname{erf} z_\sigma & z \leq z_{BL}, \\ [\alpha_+ + \alpha_- \operatorname{erf} z_\sigma] \exp z_H & z > z_{BL}, \end{cases} \quad (4.2.3)$$

where $z_H = H_{FT}^{-1}(z_{BL} - z)$. This describes the extinction profile as a function of five variables.

In the form shown, the aerosol optical thickness tends towards a constant value with height. Eventually, therefore, the impact of extinction and optical efficiency is identical within (4.2.1). Further information would be required to distinguish between these degenerate forward modelled states. One means of doing so is to measure the aerosol optical thickness with a sunphotometer. That constrains the magnitude of the extinction, such that the model only adjusts optical efficiency to fit the measurements.

There are two means by which such information could be implemented. Ideally, the χ_∞ measurement would be part of the measurement vector along with the Raman channel's profile and the four parameters of α would be retrieved within the state vector. However, in Povey et al. [2012], χ_∞ was taken as a parameter of the forward model as it was found that the retrieval neglected it in favour of φ^{ra} unless χ_∞ was measured with unrealistic accuracy.

Further, the published work used a simpler model of extinction in which χ_∞ could be easily applied as a parameter,

$$\alpha(z) = \begin{cases} \frac{\chi_\infty}{z_{BL} + H_{FT}} & z \leq z_{BL}, \\ \frac{\chi_\infty}{z_{BL} + H_{FT}} \exp z_H & z > z_{BL}. \end{cases} \quad (4.2.4)$$

That model was found to be sufficient in most circumstances, but this chapter will go beyond that work to explore the merits of applying either model.

As the forward model actually considers optical thickness, the exponential decay model (4.2.4) gives,

$$\chi(z) = \begin{cases} \frac{\chi_\infty}{z_{BL} + H_{FT}} z & z \leq z_{BL}, \\ \frac{\chi_\infty}{z_{BL} + H_{FT}} [z_{BL} + H_{FT} (1 - \exp z_H)] & z > z_{BL}, \end{cases} \quad (4.2.5)$$

whilst the error function model (4.2.3) gives,

$$\chi(z) = \begin{cases} \alpha_+ z + \alpha_- \sigma_{BL} \left\{ \frac{z_{BL}}{\sigma_{BL}} \operatorname{erf} \left(\frac{z_{BL}}{\sigma_{BL}} \right) - z_\sigma \operatorname{erf} z_\sigma \right. \\ \quad \left. + \pi^{-1/2} \left[\exp \left(-\frac{z_{BL}^2}{\sigma_{BL}^2} \right) - \exp(-z_\sigma^2) \right] \right\} & z \leq z_{BL}, \\ \alpha_+ [z_{BL} + H_{FT} (1 - \exp z_H)] \\ \quad + \alpha_- \left\{ H_{FT} \exp \left(\frac{\sigma_{BL}}{2H_{FT}} \right)^2 \left[\operatorname{erf} \left(z_\sigma - \frac{\sigma_{BL}}{2H_{FT}} \right) + \operatorname{erf} \left(\frac{\sigma_{BL}}{2H_{FT}} \right) \right] \right. \\ \quad \left. - H_{FT} \exp(z_H) \operatorname{erf}(z_\sigma) + z_{BL} \operatorname{erf} \left(\frac{z_{BL}}{\sigma_{BL}} \right) \right. \\ \quad \left. + \frac{\sigma_{BL}}{\sqrt{\pi}} \left[\exp \left(-\frac{z_{BL}^2}{\sigma_{BL}^2} \right) - 1 \right] \right\} & z > z_{BL}. \end{cases} \quad (4.2.6)$$

4.2.2 Overlap function

The analytic formulation of Halldórsson and Langerholc [1978] will be used to parametrize the integral of (4.2.1), though some alterations to its published form




Figure 4.2: A schematic of the optical axes of a lidar system.

were required, outlined as they occur through the following derivation (which is adapted from chapter 7.4 of Measures [1992]). This model provides a range of degrees of freedom and has been used successfully as the basis of other techniques [e.g. Mao et al., 2012]. Stelmaszczyk et al. [2005] and Biavati et al. [2011] in particular considered a variation that neglects the telescope's secondary mirror. Though the desire to simplify the problem can be appreciated, its omission from a monostatic system obfuscates the definition of the range of minimum overlap their analyses require and was shown to produce inconsistent overlap functions in Povey et al. [2012].

Consider a cylindrical coordinate system (R, μ, Θ) , where R is the distance of the scatterer from the telescope, μ is its radial distance from the telescope's optical axis, and Θ is the angular coordinate in that surface (zero along the perpendicular bisector of the axes), as shown in figure 4.2. The effective area is then defined as,

$$A(R, \epsilon) \equiv A_0 O(R, \epsilon) \quad (4.2.7)$$

$$= \frac{r_T^2}{w^2} \int_0^\infty \int_0^{2\pi} \psi(R, \mu, \Theta) \xi(R, \mu, \Theta) \mu \, d\mu \, d\Theta, \quad (4.2.8)$$

where r_T is the radius of the telescope's primary mirror; w is a characteristic radius of the laser beam as a function¹ of R ; and any wavelength dependence is neglected.

By neglecting wavelength, the only function of λ remaining in the calibration function is the optical efficiency, which would imply its shape is constant across all

¹In this section, the R dependence of variables will be omitted from equations for brevity.





Figure 4.3: A Cassegrain telescope modelled by a lens, radius r_T and focal length f , with a central obstruction, radius r_o . An object a distance R from the lens and radial distance μ from the optical axis produces an image a distance R' beyond the focal plane. A detector, (highly exaggerated) radius r_a , is positioned a distance Δ beyond the focal plane, such that it observes a spot of radius r'_T centred a distance μ' from the optical axis, with a spot of radius r'_o obscured.

instrument channels. In reality, the calibration function will vary from channel to channel due to inhomogeneities in components producing wavelength-dependent behaviour. Such effects will not be explicitly included within the forward model as they are virtually impossible to assess *in situ*, but their effect on the retrieval will be investigated.

A Cassegrain telescope, focal length f , can be modelled by the lens system of figure 4.3. The radii of the lens, central obstruction, and exit aperture are taken as the radii of the telescope's primary and secondary mirrors and the detector's aperture (r_T, r_o, r_a), respectively.

The flux that couples with the detector ξ is the fraction of the image spot that passes through the aperture. Defining the function $\mathcal{A}(r_1, r_2; \mu)$ as the overlap area

Figure 4.4: Illustration that $\mathcal{H} \mu d\mu d\Theta = d\mathcal{A}(\mu, w; d)$ where the infinitesimal area represents a ring intersecting the circular laser beam.



of two circles, radii r_1 and r_2 , with distance μ between their centres,

$$\xi(R, \mu, \Theta) = \frac{\mathcal{A}(r_a, r'_T; \mu') - \mathcal{A}(r_a, r'_o; \mu')}{\pi r_T'^2}, \quad (4.2.9)$$

where the first term represents the incident light; the second term represents the area obscured by the secondary mirror; and primed variables represent transformations into the plane of the detector. It is shown in Section A.6 that this can be rewritten as,

$$\xi(R, \mu, \Theta) = \frac{1}{\pi} \left(\frac{\gamma}{\nu r_T} \right)^2 [\mathcal{A}(a, \rho_T; \mu) - \mathcal{A}(a, \rho_o; \mu)], \quad (4.2.10)$$

where $\gamma = 1 + \Delta/f$; $\nu = |\gamma - f^{-2}\Delta R|$; $a = Rr_a/\gamma f$; and $\rho_{T,o} = \gamma^{-1}\nu r_{T,o}$.

Substituting into (4.2.8),

$$A(R, \epsilon) = \frac{1}{\pi} \left(\frac{\gamma}{\nu w} \right)^2 \int_0^\infty \int_0^{2\pi} \psi(R, \mu, \Theta) [\mathcal{A}(a, \rho_T; \mu) - \mathcal{A}(a, \rho_o; \mu)] \mu d\mu d\Theta. \quad (4.2.11)$$

At this stage, any beam profile could be assumed. These integrals would then be evaluated numerically to determine the overlap function. However, this would be an incredibly expensive calculation and, as accurate measurement of the beam profile is rather difficult, not likely to improve the quality of its result.

For simplicity, a top-hat beam profile is assumed,

$$\psi(R, \mu, \Theta) = \mathcal{H}(w - \sqrt{\mu^2 + d^2 - 2d\mu \cos \Theta}) \quad (4.2.12)$$

$$w = r_L + \phi_L R \quad (4.2.13)$$

$$d = \sqrt{(\delta + \phi_{\parallel} R)^2 + (\phi_{\perp} R)^2}, \quad (4.2.14)$$

where \mathcal{H} is the Heaviside step function; the beam width has been assumed to increase linearly from an initial radius r_L with divergence ϕ_L ; d is the separation of

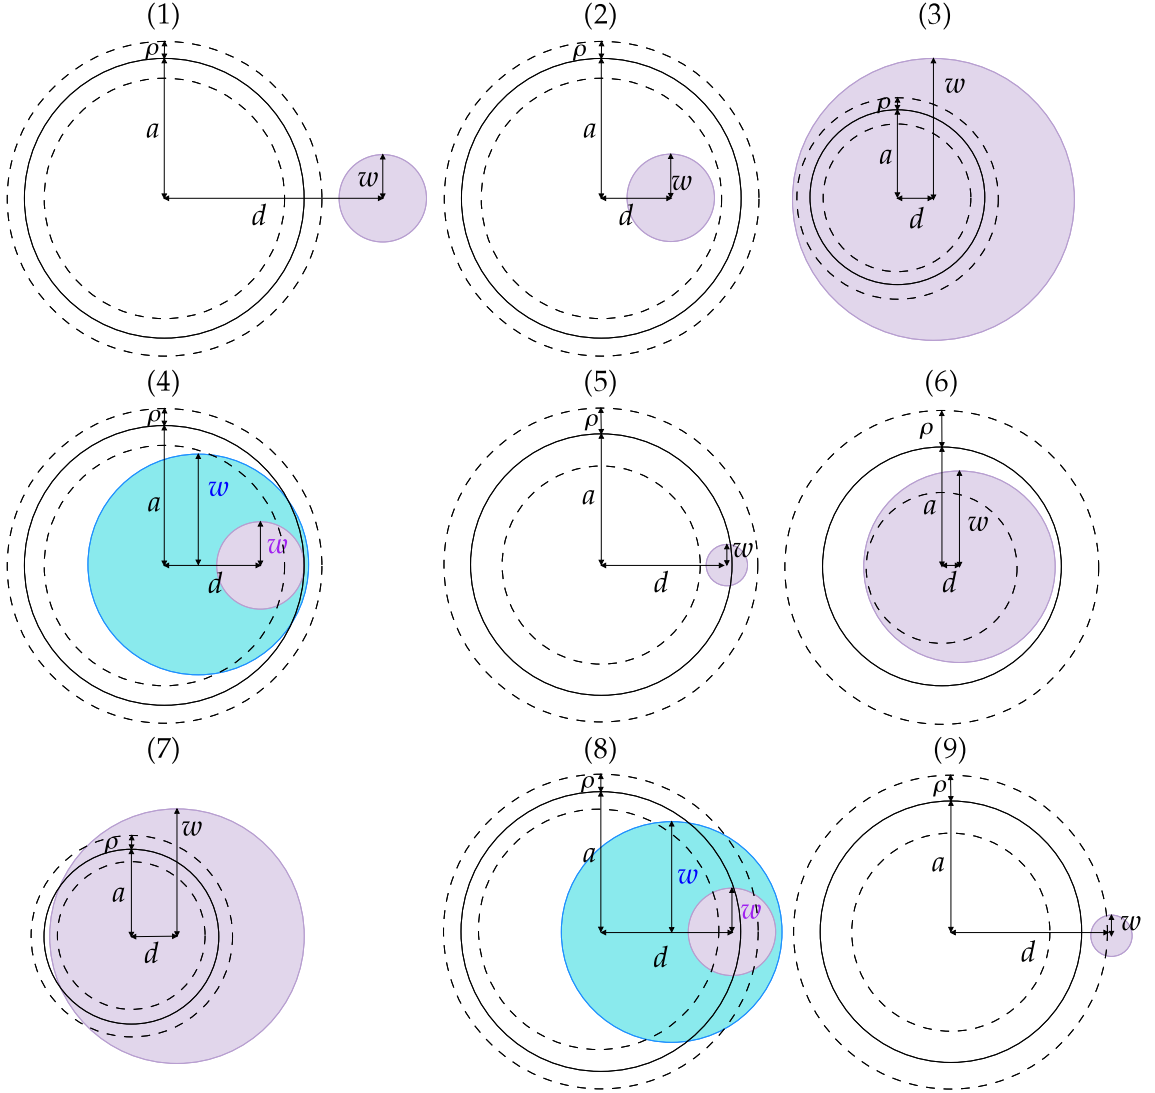


Figure 4.5: Illustration of the nine cases of S_H , viewed in the scatterer's plane for $a > \rho$. **Solid:** The telescope FOV with radius a . **Dashed:** The region $a - \rho < \mu < a + \rho$ from which light is only partially detected. **Purple:** Laser beam of radius w . **Blue:** Where necessary, a second beam with $w > d$.

the optical axes; $\phi_{\parallel, \perp}$ are the angles between those axes parallel and perpendicular to the plane $\Theta = 0$; and δ is their separation at $R = 0$. Using these, (4.2.11) can be cast into a substantially more efficient form by noting that \mathcal{H} represents a circle of radius w centred a radial distance d from the telescope's axis and $\mu d\mu d\Theta$ represents an infinitesimal ring of radius μ about that axis (see fig. 4.4). Hence, \mathcal{A} can be used to rewrite (4.2.11) as,

$$A(R, \varepsilon) = \frac{1}{\pi} \left(\frac{\gamma}{\nu w} \right)^2 \int_{\mu=0}^{w+d} [\mathcal{A}(a, \rho_T; \mu) - \mathcal{A}(a, \rho_o; \mu)] d\mathcal{A}(\mu, w; d). \quad (4.2.15)$$

The effective area can therefore be expressed as,

$$A(R, \varepsilon) = \left(\frac{\gamma}{\nu w} \right)^2 [S_H(\rho_T, a, w, d) - S_H(\rho_o, a, w, d)], \quad (4.2.16)$$

where the problem is now reduced to solving the integral,

$$S_H(\rho, a, w, d) = \frac{1}{\pi} \int_{\mu=0}^{w+d} \mathcal{A}(a, \rho; \mu) d\mathcal{A}(\mu, w; d). \quad (4.2.17)$$

There are nine possible means of combining the laser beam and detector area, sketched in figure 4.5. Numerically stable formulations of S_H for each case are derived in Section A.7.

4.2.3 Summary

The forward model for the retrieval of the calibration function from a Raman profile along a vertical path is,

$$\varphi_i^{\text{ra}} = \frac{M E_i^{\text{ra}}}{1 + \frac{\tau_d}{\tau_b} E_i^{\text{ra}}} \quad (4.2.18)$$

$$E_i^{\text{ra}} = \zeta^{\text{ra}} A(R_i, \varepsilon) N_i \frac{2}{c \tau_L} \ln \left[\frac{R_i (R_i + \frac{c}{2} \tau_b - \frac{c}{2} \tau_L)}{(R_i + \frac{c}{2} \tau_b) (R_i - \frac{c}{2} \tau_L)} \right] \times \exp \left[-(\sigma_R^{\text{el}} + \sigma_R^{\text{ra}}) \mathcal{N}_i - \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \chi_i \right] + E_B^{\text{ra}} \quad (4.2.19)$$

$$\mathcal{N}_i = \int_0^{R_i} N(R') dR' \quad (4.2.20)$$

$$\zeta^{\text{ra}} = \langle E_L \rangle \eta^{\text{ra}} \frac{\lambda^{\text{el}}}{hc} \left| \frac{d\sigma}{d\Omega} \right|_{\lambda^{\text{el}}, \lambda^{\text{ra}}}, \quad (4.2.21)$$

where ζ^{ra} contains all constant terms, such that any subsequent mention of laser energy will be quoted relative to the average during a calibration period. Two models for the aerosol optical thickness are proposed — (4.2.5) and (4.2.6).

The effective area is represented by,

$$A(R, \varepsilon) = \left(\frac{\gamma}{\nu w} \right)^2 [S_H(\rho_T, a, w, d) - S_H(\rho_o, a, w, d)], \quad (4.2.22)$$

where $\varepsilon = (r_T, r_o, r_a, f, \Delta, \delta, r_L, \phi_L, \phi_{\parallel}, \phi_{\perp})$ and,

$$\begin{aligned} \gamma &= 1 + \frac{\Delta}{f} & \nu &= \left| \gamma - \frac{\Delta}{f^2} R \right| & \rho_{T,o} &= \frac{\nu}{\gamma} r_{T,o} \\ a &= \frac{R r_a}{\gamma f} & w &= r_L + \phi_L R & d &= \sqrt{(\delta + \phi_{\parallel} R)^2 + (\phi_{\perp} R)^2}. \end{aligned}$$

To optimise performance using Rhomberg integration [Press et al., 1992], it is shown in Section A.7 that,

$$S_H(\rho, a, w, d) = \begin{cases} 0 & d - w \geq a + \rho, \\ \pi w^2 \min[a^2, \rho^2] & w + d \leq |a - \rho|, \\ \pi a^2 \rho^2 & w - d \geq a + \rho, \\ w^2 \mathbb{A}(a, \rho; w + d) & |a - \rho| < w + d < a + \rho \\ & \text{and } |w - d| \leq |a - \rho|, \\ w^2 \mathbb{A}(a, \rho; w + d) + \frac{1}{\pi} \int_{|a-\rho|}^{w+d} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu} & |a - \rho| < w + d < a + \rho \\ & \text{and } d - w > |a - \rho|, \\ w^2 \mathbb{A}(a, \rho; w + d) + \mathbb{B} & |a - \rho| < w + d < a + \rho \\ & \text{and } w - d > |a - \rho|, \\ w^2 \mathbb{A}(a, \rho; w + d) + \frac{1}{\pi} \int_{w-d}^{w+d} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu} & w + d \geq a + \rho \text{ and } \\ & |a - \rho| < w - d < a + \rho, \\ \mathbb{B} + \frac{1}{\pi} \int_{w-d}^{a+\rho} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu} & w + d \geq a + \rho \text{ and } \\ & |w - d| \leq |a - \rho|, \\ 2 \min[a^2, \rho^2] \mathbb{C}(|a - \rho|, w, d) & w + d \geq a + \rho \text{ and } \\ & |w - d| \leq |a - \rho|, \\ + \frac{2}{\pi} \int_{|a-\rho|}^{a+\rho} \mathbb{A}(a, \rho; \mu) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu & w + d \geq a + \rho \text{ and } \\ & |a - \rho| < d - w < a + \rho, \\ \frac{1}{\pi} \int_{d-w}^{a+\rho} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu} & \end{cases} \quad (4.2.23)$$

where,

$$\Upsilon(a, \rho, \mu) = 2(r_1^2 r_2^2 + r_1^2 \mu^2 + r_2^2 \mu^2) - r_1^4 - r_2^4 - \mu^4 \quad (4.2.24)$$

$$\mathbb{A}(a, \rho; \mu) = a^2 \cos^{-1} \left(\frac{\mu^2 + a^2 - \rho^2}{2\mu a} \right) + \rho^2 \cos^{-1} \left(\frac{\mu^2 + \rho^2 - a^2}{2\mu \rho} \right) - \frac{1}{2} \sqrt{\Upsilon(a, \rho, \mu)} \quad (4.2.25)$$

$$\mathbb{B} = \frac{(w - d)^2 - a^2 - \rho^2}{4} \sqrt{\Upsilon(a, \rho, w - d)} + 2a^2 \rho^2 \tan^{-1} \sqrt{\frac{(w - d)^2 - (a - \rho)^2}{(a + \rho)^2 - (w - d)^2}} \quad (4.2.26)$$

$$\begin{aligned} \mathbb{C}(x, w, d) &= \frac{x^2}{2} \cos^{-1} \left(\frac{x^2 + d^2 - w^2}{2xd} \right) - \frac{\sqrt{\Upsilon(x, w, d)}}{4} \\ &\quad - \frac{w^2}{2} \tan^{-1} \left(\frac{w^2 + d^2 - x^2}{\sqrt{\Upsilon(x, w, d)}} \right) + \frac{\pi w^2}{4}. \end{aligned} \quad (4.2.27)$$

The form of \mathbf{K} is summarised in Section A.3. The state and measurement vec-

tors are,

$$\mathbf{x} = (\delta, \phi_{\parallel}, \phi_{\perp}, r_L, \phi_L, r_T, r_o, r_a, f, \Delta, \tau_L, \eta, z_{BL}, H_{FT}, \chi_{\infty})^T \quad (4.2.28)$$

$$\text{or } (\delta, \phi_{\parallel}, \phi_{\perp}, r_L, \phi_L, r_T, r_o, r_a, f, \Delta, \tau_L, \eta, z_{BL}, \alpha_+, \alpha_-, \sigma_{BL})^T \quad (4.2.29)$$

$$\mathbf{y} = (\varphi_0^{\text{ra}}, \varphi_1^{\text{ra}}, \dots, \varphi_{m-1}^{\text{ra}}, \chi_{\infty})^T. \quad (4.2.30)$$

4.3 Expected properties and distribution

4.3.1 *A priori* assumptions

Most elements of this state vector correspond to physical quantities so their *a priori* can be based upon direct measurements. As in the previous chapter, the RACHEL system during its campaign at CFARR will be used as the basis for simulations (see table 4.1). It should be noted that the ϕ_L estimate derives from a manufacturer's quote of < 1.5 mrad transmitted through a factor-five beam expander and that τ_L was quoted as 5–7 ns.

The system is presumed well-aligned, such that parameters describing the misalignment are ascribed an *a priori* of zero with variances based on physically reasonable ranges:

- The system is coaxial, such that δ is limited to a few centimetres by the size of the reflecting mirror.
- At the very least, it can be assumed that the system will be sufficiently well-aligned that there is some sensitivity through the first few kilometres. Investigation of (4.2.22) shows this requires $|\phi_{\parallel, \perp}| \lesssim 0.4$ mrad.
- The length of the telescope's eyepiece limits Δ to a few centimetres.

A rough estimate of $\zeta = 1.5 \times 10^{-19}$ counts $\text{m}^2 \text{sr}^{-1}$ can be produced from manufacturer's specifications and some laboratory measurements for the optical components. However, as noted in the last chapter, the laser energy has been seen to vary by up to a factor of four, inspiring a very conservative estimate for the uncertainty: 6×10^{-19} counts $\text{m}^2 \text{sr}^{-1}$.

The OPAC model [Hess et al., 1998] gives a scale height of 2 km over which aerosol abundance decays. Considering different types of aerosol, Seinfeld and Pandis [1998] gave scale heights of 0.9 – 2 km (excluding polar aerosols). These values are consistent with those observed by backscatter sondes (figure 2.5), such that $H_{FT} = 2(1)$ km is deemed a reasonable estimate. However, when using the

exp extinction model, the exponential simulates the sharp decrease in extinction across the entrainment layer, such that the value should be more representative of the entrainment layer thickness. Following Povey et al. [2012], 200(100) m will be adopted.

The remaining variables for the extinction profile are more difficult to constrain. The variation of PBL height remains an area of intense study, showing significant diurnal, seasonal, yearly, and regional differences [e.g. Seibert et al., 2000; Brooks, 2003; Wiegner et al., 2006; He et al., 2008; Barlow et al., 2011; Granados-Munoz et al., 2012]. However, for northern Europe, 1 km is a suitable order of magnitude. Similarly, the entrainment layer thickness is of order 100 m.

The aerosol optical thickness at CFARR has been monitored for several years by an AERONET station. Figure 4.6 shows a histogram of the level 2.0 aerosol optical thickness at 380 nm observed during 2006–10. Superimposed on this are rough fits of a normal and log-normal distribution with means 0.14 and -1.7 and widths 0.28 and 0.6, which shall be used as *a priori* estimates of χ_∞ .

If extinction is neglected in the free troposphere, $\alpha_+ \simeq \alpha_-$ and,

$$\chi_\infty = \alpha_+[z_{BL} + H_{FT}] + \alpha_- \left\{ H_{FT} \exp \left(\frac{\sigma_{BL}}{2H_{FT}} \right)^2 \left[\operatorname{erf} \left(\frac{\sigma_{BL}}{2H_{FT}} \right) - 1 \right] + z_{BL} \operatorname{erf} \left(\frac{z_{BL}}{\sigma_{BL}} \right) + \frac{\sigma_{BL}}{\sqrt{\pi}} \left[\exp \left(-\frac{z_{BL}^2}{\sigma_{BL}^2} \right) - 1 \right] \right\} \quad (4.3.1)$$

$$\simeq 3 \times 10^3 \alpha_+ - 1 \times 10^3 \alpha_- \quad (4.3.2)$$

$$\therefore \alpha_\pm \simeq 5 \times 10^{-4} \chi_\infty. \quad (4.3.3)$$


The uncertainty of this should be greater than what simple arithmetic would imply to account for variations in α_{FT} .

4.3.2 Distribution of state variables

One of the assumptions of optimal estimation is that the PDF of all variables is normally distributed or at the very least symmetric. One means to sample the PDF of a forward model is the stationary Metropolis-Hastings algorithm [Chib and Greenberg, 1995]. This can be summarised as:

1. Generate a random state vector, normally distributed about the first guess of the retrieval with variance given by the *a priori* covariance matrix.
2. For that state, evaluate the forward model and its cost \mathcal{C} with respect to a simulated profile.

Figure 4.6: **Black:** Distribution of $\chi_\infty(380\text{ nm})$ observed by the CFARR AERONET station from Feb 2006 to Sep 2010. **Blue:** Linear *a priori* distribution. **Red:** Log-normal *a priori* distribution.



3. Determine the probability that the state is a better description than the first guess,

$$P(\mathbf{x}) = \exp \left\{ \left(\frac{m}{2} - 1 \right) [\ln \mathcal{C}(\mathbf{x}) - \ln \mathcal{C}(\mathbf{x}_a)] - \frac{\mathcal{C}(\mathbf{x}) - \mathcal{C}(\mathbf{x}_a)}{2} \right\}. \quad (4.3.4)$$

4. Generate a random number, uniformly distributed between 0 and 1. If it is less than $P(\mathbf{x})$, retain the state as a member of the PDF.
5. Repeat (here 10^4 iterations). The PDF may be approximated by the distribution of retained states.

The algorithm should be tuned such that the rate of retention is $\sim 25\%$. This is done by multiplying the covariance matrix by a constant (here found to be 0.08 by trial and error), which is a simple if not completely rigorous methodology.

As shown in figures 4.7 and 4.8, most of the PDFs are approximately Gaussian with the exception of:

- Δ , which has a large positive skew;
- δ , which would have a symmetric distribution about zero but cannot be negative as it represents a length; and
- ζ , which is slightly negatively skewed.




Figure 4.7: State vector PDFs as estimated by the stationary Metropolis-Hastings algorithm for the erf configuration. Dashed lines show the value of x_0 . For the ϕ plot, blue is ϕ_{\parallel} and red ϕ_{\perp} .




Figure 4.8: As fig. 4.7, but for the exp configuration. The cause of the central peaks is not known.




Figure 4.9: The effect of Δ on the effective area for an otherwise well-aligned RACHEL system. Note how positive and negative values produce a different evolution.

A negative value of δ could be interpreted as reversing the sign of ϕ_{\parallel} , such that it may be permitted to take negative values. Also, in the latter two cases, if the logarithm of the variable is retrieved instead, its distribution is symmetric (see figs. B.7 and B.8). However, that trick cannot be applied to Δ , as it can be negative.

Its asymmetric distribution is physically sensible (see fig. 4.9). As Δ is increased from a large, negative value, the magnitude of the calibration function increases at all levels, tending towards the properly-aligned solution. Increasing it above zero improves sensitivity near the surface whilst decreasing it at height. Positive and negative values clearly have different Jacobians and so cannot be expected to produce a symmetric PDF. This would be expected to cause the retrieval to unduly favour negative Δ states and incorrectly estimate its error contribution.

The exp configuration also shows a slight skew in the distribution of H_{FT} . In Povey et al. [2012], its logarithm was retrieved both to correct that skew and to prevent the retrieval of negative scale heights, which is otherwise common as the *a priori* distribution overlaps significantly with negative values. At that time, the other variables were observed to be symmetrically distributed, which likely implies it considered an improper range of states to form its PDFs.

The retrieval will consider the same range as that of Chapter 2: 0.1–5 km. The first guess is taken from the published algorithm, which was found to minimise the number of iterations required for solution and avoids the null Jacobian given by $\mathbf{x} = 0$ for some variables.

Table 4.1: *A priori* state vector, with standard deviations in brackets, for the calibration function retrieval. Superscripts indicate linear and logarithmic retrievals.

Variable	Value	Variable	Value	Variable	Value
Δ	0.0(1) m	ϕ_{\parallel}	0.0(4) mrad	ϕ_{\perp}	0.0(4) mrad
$2r_L$	35(5) mm	ϕ_L	0.3(1) mrad	f	2.00(1) m
$2r_T$	203(1) mm	$2r_o$	75(10) mm	$2r_a$	0.40(4) mm
τ_L	6(1) ns	δ^{lin}	0.0(1) m	ζ^{lin}	$1.5(60) \times 10^{-19} \text{ m}^2 \text{ sr}^{-1}$
z_{BL}	1(1) km	δ^{log}	-4.6(80)	ζ^{log}	-43.4(30)
$\chi_{\infty}^{\text{lin}}$	0.14(28)	$(H_{FT}^{\text{exp}})^{\text{lin}}$	0.2(1) km	$(H_{FT}^{\text{erf}})^{\text{lin}}$	2(1) km
$\chi_{\infty}^{\text{log}}$	-1.7(6)	$(H_{FT}^{\text{exp}})^{\text{log}}$	5.3(20)	$(H_{FT}^{\text{erf}})^{\text{log}}$	7.6(20)
σ_{BL}	0.1(1) km	$\alpha_{\pm}^{\text{lin}}$	70(150) Mm^{-1}	$\alpha_{\pm}^{\text{log}}$	-8.6(6)

As published, this retrieval used $\Gamma_0 = 10^2$. It was found necessary to increase this to 10^5 . In the retrieval, Γ_i can be thought of as the distance between successive states and drives the retrieval between a Gauss-Newton and greatest descent minimisation procedure through the inverted matrix of (1.5.10). As the Gauss-Newton method is most applicable to a smoothly curved surface of the cost function, it is desirable to start with a large Γ_0 (the greatest descent method) until sufficiently close to the minimum. Hence, when the measurement error is smaller, it is necessary to use a larger Γ_0 , though its exact value should be unimportant to the retrieval. By trial-and-error sampling, it was clear that the solution for 10^5 was more similar to those for 10^4 and 10^6 than any other order was to its neighbours².

4.3.3 Simulated data

Simulation of data is straightforward with this forward model, though it may eventually be desirable to implement a different representation of the detector nonlinearity. Seven representative profiles are produced to assess the performance of the various possible forward model configurations. The same standard atmosphere will be used as before (table B.3). The profiles represent (see table B.4):

1. A reasonably well-aligned system with all easily measured parameters taking their known values (table 1.1) observing a PBL of typical optical thickness (0.18) at night over 60,000 laser shots.
2. As case 1, but perturbing the easily measured parameters.

²The methods of L-curve analysis [Hansen, 1992; Pounder et al., 2012] have been applied to determine constants such as this in other studies but do not appear to be valid for this retrieval.

3. As case 1, but observing an optically-thicker PBL (0.37).
4. As case 1, but for a less well-aligned system.
5. A different alignment observing an optically thin PBL (0.06) with all easily measured parameters taking their known values.
6. Yet another alignment observing a somewhat thicker PBL (0.11) with the easily measured parameters perturbed slightly from their known values.
7. A poorly-aligned system where the calibration function peaks at 400 m.

The χ_∞ measurement is determined analytically from the simulated state vector. An error of 2 % is then assumed, in line with practical ability, though noise is not added.

4.4 Configuration assessment

Though the state vector PDFs indicate that some variables should be retrieved logarithmically, this does not comment on the practical implications of doing so. Further, many variables must be positive, for which retrieving the logarithm may be beneficial. As shown in Chapter 2, the configuration of the retrieval can impact its behaviour in a way which is not intuitively obvious by changing which areas of state space are considered more viable solutions.

Though variables such as r_T and f must be positive, their *a priori* distribution is sufficiently constrained that the probability of exploring negative values should be negligible, though this is not guaranteed³. An ensemble of simulated profiles were produced from randomised \mathbf{x} and $r_T, r_a, r_o, r_L, \phi_L, f, z_{BL}$, and σ_{BL} were never observed to approach negative values. Hence, these shall be retrieved linearly with a lower limit of zero applied. As they can be negative, there is also no reason to retrieve $\phi_{\parallel}, \phi_{\perp}, \Delta$, and τ_b by any means other than linearly. The aerosol parameters α_{\pm} and χ_∞ must be positive but were sometimes retrieved to be negative. Thus, the behaviour of the retrieval will be investigated for both the exp and erf aerosol models over the 16 possible permutations of linear or logarithmic retrieval of $\delta, \zeta, H_{FT}, \alpha_{\pm}$, and χ_∞ .

A range of results are returned by the various configurations, a sample of which are highlighted in figure 4.10. The retrievals are accurately capturing the shape of

³The parameter values will not necessarily correspond to physically meaningful quantities due to degeneracy in $A(R)$ producing equivalent solutions from very different inputs.




Figure 4.10: Components of C retrieved from simulated data, showing three configurations deemed representative of the behaviour of all 32. The central plot shows C normalised by its maximum to highlight its functional form over its magnitude.




Figure 4.11: Retrieved residuals ($y - \mathbf{F}$) from noise-free data for the configurations shown in fig. 4.10.

the calibration function, as shown in the middle column, but have more difficulty in establishing its magnitude, as shown by the greater spread of the top row. The occasions where the retrieval of the shape fails (such as the green line of case 3) correspond to where the retrieved profile of optical thickness, as shown on the bottom row, is noticeably wrong. The more poorly aligned case 7 is less well retrieved than the others whilst the magnitude of χ_∞ appears to be unimportant as retrievals from both cases 3 and 5 present greater variations than case 1.

All retrievals return a solution with $\mathcal{C} < 1.2$ and the vast majority converge in fewer than 30 iterations. To investigate any biases in the forward model, these retrievals were repeated with noise-free data and the difference between the simulated and retrieved lidar profiles is shown in figure 4.11. These are found to be minimal in all but case 7. Figure 4.12 explores this difficult case by considering the change in the retrieved \mathcal{C} as the SNR (i.e. integration time) is increased. Though the results never reproduce the simulated profile, they do tend towards it as the SNR increases, indicating that the issue is not with the retrieval's convergence.

Comparing the various configurations in fig. 4.10, the *exp* modes behave very similarly — slightly overestimating \mathcal{C} at the top of the PBL due to an inability to fit a large value of α_{FT} (e.g. case 3). These are worse than *erf* modes in all but case 7. The *erf* modes retrieving δ linearly tend to overestimate case 5 and logarithmic retrievals instead overestimate case 6 (not shown). Both configurations tend to retrieve a decrease in \mathcal{C} with height for case 3 (though they remain consistent with the truth).

However, this has only been a qualitative consideration of the configuration's




Figure 4.12: Retrieved calibration function for case 7 as a function of the SNR. Line style indicates the retrieval's configuration while colour denotes the averaging period (a proxy of SNR).





Figure 4.13: The performance of various calibration function configurations. **Top:** RMSD of the retrieved profile from truth. **Bottom:** One minus the correlation between the retrieved and true profiles. **Left:** Calibration function. **Right:** Aerosol optical thickness. **Colour:** The number of iterations required to converge. **Diamonds:** erf configuration retrievals. **Crosses:** exp configuration retrievals. **Displacement:** Retrievals of $\ln \delta$ are displaced to the right of the axis; linear δ retrieval are displaced to the left.

performance. Figure 4.13 presents, for all proposed configurations, the RMSD between the retrieved and true profiles and their correlation. Accurate retrievals converge in around 20 iterations, though cases 6 and 7 generally required more.

Figure 4.14: Overlap function for the CUV as determined by CFARR by direct inversion of (1.3.17) for observations on 21 Sep 2005.



The erf configurations are always more consistent with the truth than the exp configuration for cases 1–4, with the opposite being true for cases 5–7 (though to a lesser magnitude). The mode of δ tends to produce better retrievals, but not consistently and is mostly presented to disperse the data points.

The RMSD were then used to rate each configuration. All deviations were normalised by the smallest encountered for each simulated profile. These ‘scores’ were then averaged to find the configurations that most frequently produced the smallest deviation.

1. Considering cases 1–4, the best two configurations used the erf model with $\ln \zeta$ and linear δ & H_{FT} .
2. Considering cases 5–7, the best two used the exp model with either $\ln \delta$ and $\ln \zeta$ or a logarithmic retrieval of all four variables.
3. Across all seven cases, the best two configurations were erf mode with $\ln \delta$, $\ln \alpha_{\pm}$, and linear H_{FT} . It is interesting to note that the third best choice here was the best choice for cases 1–4.

4.5 Sensitivity

The accuracy of these six ‘best’ configurations was further assessed by considering additional simulated profiles representing poorly aligned circumstances:

8. As case 1, but with $\Delta = 18$ mm.
9. As case 1, but with $\phi_{\parallel} = 0.37$ mrad.

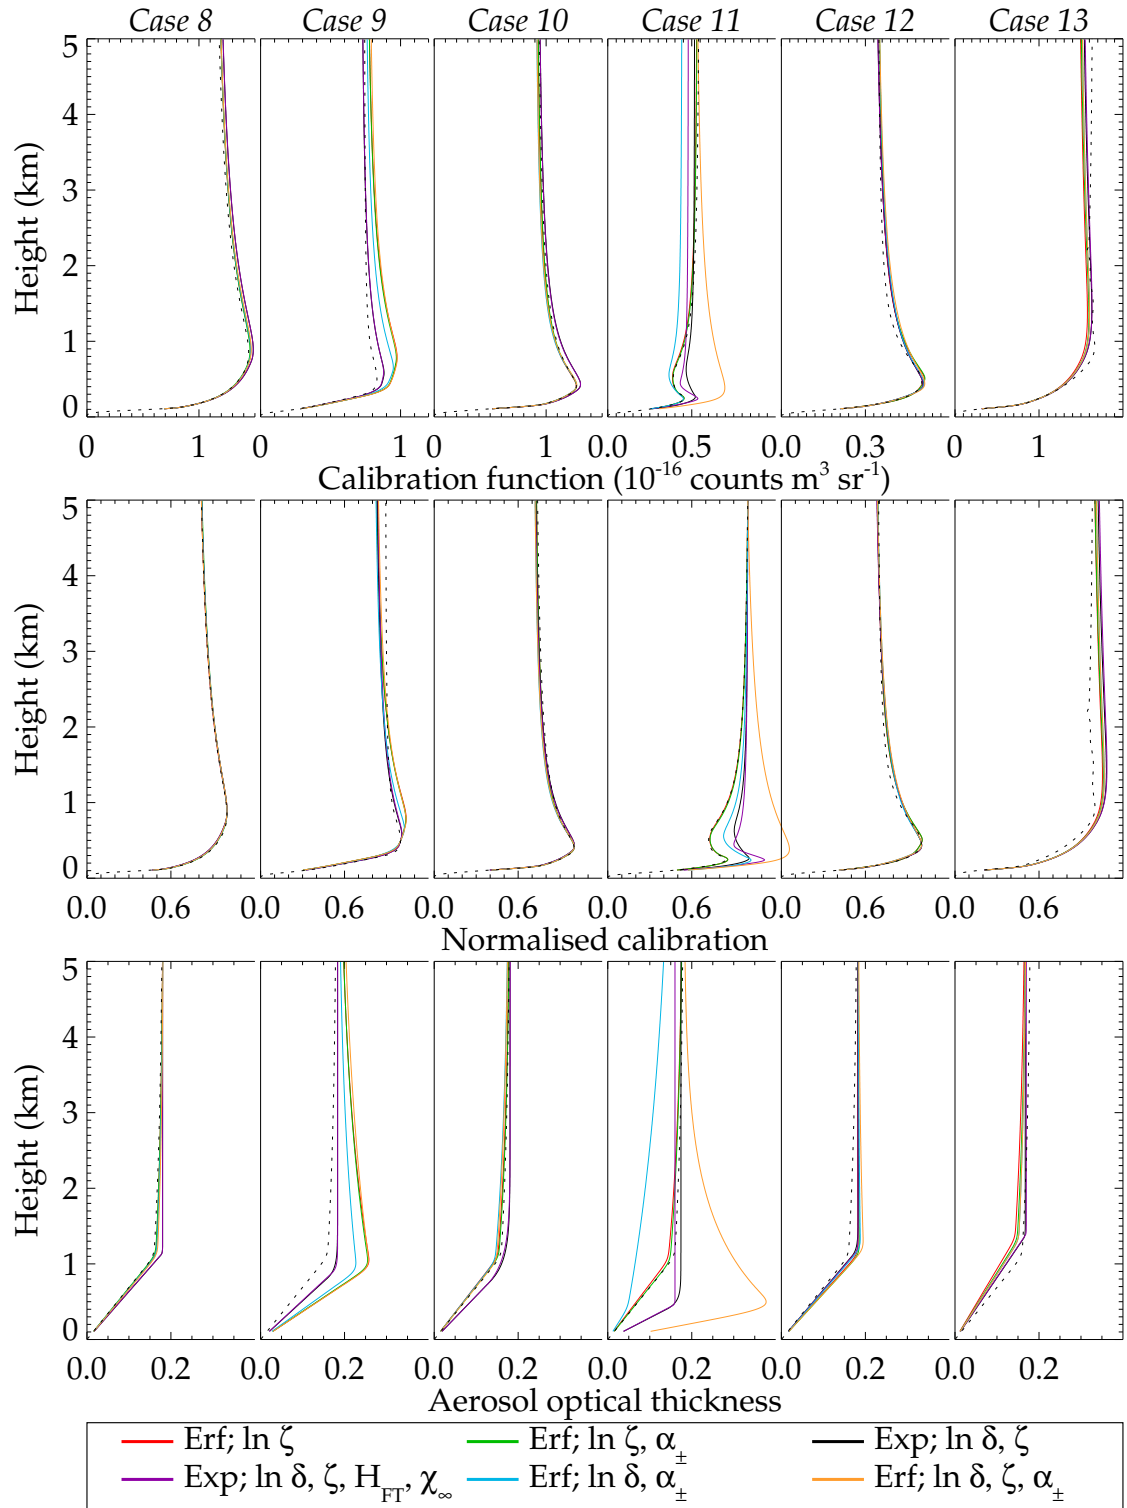



Figure 4.15: As fig. 4.12, but for the six ‘best’ configurations of Section 4.4 and considering cases 8–13.

Figure 4.16: Perturbed aerosol profiles used for data simulation. **Purple:** Idealised aerosol layers. **Blue:** A thin, idealised cloud. **Green:** Random, multiplicative noise. **Red:** A linear decrease in extinction.



10. As case 1, but with $\phi_{\parallel} = 0.17$ mrad and $\Delta = 18$ mm.
11. As case 1, but with $\phi_{\parallel} = 0.37$ mrad and $\phi_{\perp} = 0.21$ mrad.
12. As case 1, but with $\phi_{\parallel} = 0.37$ mrad, $\phi_{\perp} = 0.21$ mrad and $\Delta = 0.15$ mm.

A thirteenth profile was also prepared to mimic the impact of optical components not included in the forward model. The overlap functions for each channel of the CUV have been measured by direct inversion of (1.3.17) in clear conditions (fig. 4.14). To remove noise, a tenth-order polynomial was fit to the ratio of these curves over the first 2.5 km. That polynomial is then used to perturb the overlap function from case 1 to mimic the effect of a real optical system.

The results are shown in figure 4.15. These are clearly not as accurate as in fig. 4.10 but still present a useful approximation. No one configuration is clearly superior. The very poorly aligned case 11 is an outlier producing a wide variety of answers. This behaviour is driven by subtle differences between the definitions of the *a priori* in each configuration such that the instances producing a good retrieval are potentially coincidental. However, the poor retrievals can be readily identified by their large cost (> 2) and a calibration of this form would indicate significant practical difficulties, so should never be observed. The retrieval is also not overly successful in case 13, which will be considered further shortly.

Though the retrieval considers both the calibration function and the aerosol profile, only the calibration function is of further interest. Thus, it is worth considering how its solution evolves as the aerosol profile is altered. The parameters for case 1 were used to produce additional simulations with various alterations made to the aerosol extinction profile. Four examples were considered where the shape of the profile is maintained, but its magnitude is scaled to produce χ_{∞} between

0.05 and 0.4. Four further examples, shown in figure 4.16, add features not included within the forward model to the extinction profile:

- Idealised aerosol layers of the form outlined in (2.5.1);
- A thin cloud modelled by a Gaussian;
- Random, multiplicative noise of standard deviation 2 %; and
- A linear decrease in extinction from a maximum chosen to give an optical depth of 0.18.

The impact of the aerosol perturbations on the retrieval is small. To emphasise them, the difference between the retrieved and simulated profile is presented instead of the profiles directly. Shown in figure 4.17, they are at most 5–10 % and are larger for the exp configurations. From the top row, it is evident that as χ is increased the retrieval increasingly diverges from the true profile, possibly because the absolute error on χ_∞ increases (leaving ζ less well-constrained). Indeed, it is not sensible to attempt calibration when $\chi > 0.2$ as the extinction profile is unlikely to be well modelled by the simple forms assumed.

The impact of the features of fig. 4.16 is larger, as shown on the bottom row of fig. 4.17. Aerosol layers or a thin cloud cause an overestimation of C due to an overestimation of χ through them. Linearly decreasing extinction causes the opposite effect. However, such features should be easily identified in the ABC. By selecting data lacking such features, errors due to the aerosol profile can hopefully be constrained to those demonstrated by the multiplicative noise — a few percent.

Another consideration is the assumption that the laser has a top-hat beam shape. For the pulsed lasers used in lidar, beam energy is broadly similar to a top-hat in that it decreases sharply at its edges but will likely have an irregular shape. As the beam propagates, it tends towards a Gaussian shape which is accentuated by vignetting in the observation system. As such, even if the beam shape were measured in a laboratory, it will not necessarily bear any resemblance to that in practice. As the retrieval models one extreme of behaviour, its performance at the other should be inspected. As shown in figure 4.18, retrievals on data simulated assuming a Gaussian beam shape are discouraging. They overestimate C and χ by up to 20 % and do so in solutions with reasonable cost. This is partly because they utilise unrealistic χ profiles for which $\alpha_- > \alpha_+$ have been retrieved.

Effectively, the retrieval is using a peaked optical thickness profile to compensate for decreasing C above 2 km (caused by large values of $\phi_{\parallel,\perp}$). This is




Figure 4.17: Deviation of retrieved C from the truth when the aerosol profile has been perturbed. Note that the scale is an order of magnitude smaller than in previous figures. **Top:** The optical thickness of the profile is scaled to values on the range 0.05–0.4. **Bottom:** Features are added to the extinction profile, as shown in fig. 4.16.

clearly unphysical and undesirable, but appears to be incredibly difficult to prevent. These minima are sufficiently appealing⁴ that placing hard limits on x to force physical answers does not change the solution; it merely increases its cost. The forward model would need to be cast in a different manner — perhaps expressing α_+ as some multiple of α_- — to completely eliminate these solutions from state space. However, such contortions degrade the ability of the retrieval to converge by flattening state space. A solution to this will be introduced at the end of this section.

The sensitivity of these configurations to the retrieval’s first guess, x_0 , was investigated by producing an ensemble of 50 randomized guesses for δ , ϕ_{\parallel} , and Δ

⁴In the sense that the minima are steep and deep in state space.




Figure 4.18: Deviation of retrievals from cases 1–7 simulated with a Gaussian beam shape. Line style indicates configuration.

with variances 5 mm, 0.1 mrad, and 5 mm, respectively⁵. Shown in figure 4.19, it is clear that the erf $\ln \zeta$ configurations (top) are producing the most robust solutions and that the exp mode retrievals (middle) demonstrate excessive sensitivity. The remaining erf configurations (bottom) are also overly sensitive in the more poorly-aligned cases 3, 6, and 7, frequently converging to high-cost solutions. Hence, these four configurations will be discounted from further consideration.

Though the magnitude of Γ_0 is important, its exact value should have no effect upon the retrieved profile. Figure 4.20 shows that this is only true when using an error function aerosol profile retrieving all variables linearly except $\ln \zeta$. For this configuration, the choice of maximum retrieved height is unimportant provided it captures enough useful data, with the previously applied 5 km being sufficient (fig. B.9). The minimum height is more important than in Chapter 2’s retrievals as the calibration function varies most rapidly near the surface. Figure 4.21 shows that there are significant biases in the first few retrieved bins. It implies that if

⁵These guesses were found to have the greatest impact on the retrieval’s result. The other variables tend to only affect the retrieval when they are greatly different from their true value and then only by preventing convergence.




Figure 4.19: Deviation of retrievals from cases 1–7 using a randomised first guess for δ , $\phi_{||}$, and Δ . Colour denotes the deviation of that guess from zero, weighted by the *a priori* covariance.




Figure 4.20: Deviation of retrieval from cases 1–7 as a function of Γ_0 . Line style indicates configuration.




Figure 4.21: Deviation of retrieval from cases 1–7 as a function of minimum retrieved height. Line style indicates configuration.

the applied minimum height of 100 m is decreased, the performance with poorly aligned profiles would improve. Reanalysis of cases 7–12 with this lower bound did not significantly improve the results and so will be neglected as, in practice, these bins tend to be contaminated by detector ringing after the laser shot.

However, it is worth considering why the solutions shown in figure 4.21 are more consistent with the true profile than shown earlier in figure 4.10. This requires a closer examination of the top row of fig. 4.19. There, the y-axis has purposefully not been drawn to highlight that, for a few choices of x_0 , the retrieval converges to a *more* accurate solution than previously observed with $\mathcal{C} < 1.05$ compared to 1.1–1.2 previously. This relates to the presence of local minima of the cost function.

The state space of this retrieval can be visualised as a potholed road — mostly flat but with some dips and cracks, one of which is deeper than the others (hopefully corresponding to the correct answer). If the retrieval falls into any of the shallow dips, it tends to remain there and converge to an incorrect solution with a large cost. In most of the explorations presented, the retrieval was able to correctly identify the deepest ‘hole’ immediately and converge upon the correct answer. When the data contain features not considered within the forward model, the retrieval is less able to judge the correct ‘dip’ and so must start ‘nearer’ to it in state space to be successful. As this is not known *a priori*, a suitable response is to select a random set of starting points and run a retrieval from each of them. Of those, the solution with the lowest cost is the accepted answer. This is similar to the technique of simulated annealing though, rather than add a random perturbation to the state vector during the retrieval, it merely starts from different points.

The application of such annealing is shown in figure 4.22. The annealed solutions (solid) are mostly more accurate than retrievals using the static first guess (dashed). However, case 13 is never well-retrieved with all solutions expressing a large cost. The retrieval simply cannot capture the small-scale structure of that simulation, though the annealed solution is at least more consistent with the true profile. It is worth asking if this was a reasonable means to simulate the impact of the optical system? The profiles of fig. 4.14 are derived for the CUV whilst the simulations were modelled on RACHEL, so it may be that a discrepancy of this magnitude is simply not possible in RACHEL. Those calibrations are also, with all due respect to CFARR, excessively noisy. The published version of this retrieval had less difficulty in this matter (its fig. 6), but the perturbation considered there was a lower-order polynomial based on a qualitative estimate of the difference




Figure 4.22: The impact of annealing the first guess of the retrieval for cases 7, 13, and the Gaussian simulations of fig. 4.18. **Top:** Deviation of C retrieved from randomised x_0 , with colour indicating the cost of the solution. The solution with minimal cost is plotted in black. The solution from the static first guess shown dashed. **Bottom:** Retrieved aerosol optical thickness with and without annealing. The true profile is dotted.

between the overlaps rather than an actual measurement. This matter will receive further consideration when this retrieval is applied to actual data in Chapter 5.

The annealing also returns a worse solution for the first Gaussian simulation. The retrieved χ , shown on the bottom of fig. 4.22, is clearly unphysical. If that solution is discounted, the next best solution is no worse than non-annealed result. The same can be done for Gaussian case 5, where the original solution has the lowest cost but a more accurate solution was found for a profile using a physically sensible χ . In the majority of cases, however, the annealing has improved the retrieved solution such that, in the circumstances containing features not explicitly included within the forward model, its discrepancy from the true profile is only a few percent (a similar order to those caused by aerosol perturbations). These errors are

correlated through the profile and so are more systematic than random.

By randomising the retrieval's first guess, it is allowed to explore state space more broadly than otherwise permitted by the Gaussian *a priori* constraints. For data containing features not represented in the forward model, it permits the algorithm to find equivalent solutions that, though wrong in the strict sense of using parameter values with no physical basis, are functionally useful. This ability to fit a wide range of functions is the reason the full model of Halldórsson and Langerholc [1978] was applied rather than a simplified version — its complexity permits solutions beyond ideal circumstances. Annealing isn't a perfect solution to the challenges facing this retrieval, being very computationally expensive and not necessarily well-suited to optimal estimation, but it does give useful answers.

4.6 Errors

4.6.1 Retrieval error

As the retrieved parameters are not of particular interest, the derived S_x must be mapped onto the calibration function by,

$$S_C = K_C^T S'_x K_C, \quad (4.6.1)$$

where $K_C = \frac{dC}{dx}$, which can be derived from the expression for K in Section A.3 and S'_x is a truncated version of S_x with the rows and columns describing α removed.

Ignoring the annealing of x_0 for the moment, figure 4.23 plots the solutions for cases 1–7 with retrieved errors. These are not consistent with the truth in cases 3, 4, 6, and 7. As a simple means to assess if the errors are actually representative of the retrieval's behaviour, the noise on each measurement was re-simulated 35 times and retrievals from those are plotted in green. It is clear that the error bars are a reasonable description of the retrieval's response to measurement noise in all but case 7. It is also clear that the solutions in cases 3, 6, and 7 are outliers on the distribution of retrieved states, again demonstrating the utility of randomising the first guess.

Figure 4.24 compares, for all simulations, the solution retrieved with and without annealing to highlight how, where the retrieval was performing well, annealing does no worse but it greatly improves poor retrievals such as case 7. The figure also compares the retrievals to the results of existing techniques applied to the same




Figure 4.23: Exploration of the error derived for C , plotting the non-annealed solution in black with retrieved error in blue for cases 1–7 (left to right). This is compared to the retrievals from 35 profiles simulated with different noise (green).

data. A simple inversion of (1.3.19) is plotted in blue⁶. The iterative Wandinger correction (4.1.2) is plotted in green⁷. Both of these solutions are dominated by measurement noise, though the iterative scheme is fairly accurate (if imprecise) above 500 m. Increasing the integration time of the measurement a thousand times (equivalent to a month of observation, shown in red) eliminates the noise and most of the bias⁸. Even with this impossible improvement in the SNR, the annealed optimal estimation retrieval performs at least as well as the existing technique in all but case 13.

Figure 4.25 briefly explores the mechanics of the retrieval. The averaging kernels make it plainly obvious why Povey et al. [2012] could succeed with a truncated state vector. The retrieval is insensitive to most of variables (the diagonal would be unity for a perfect retrieval) and so has insufficient information to deviate from the *a priori*. That is as expected and, if efficiency of the retrieval were a concern, those elements could be removed from the state vector and considered parameters of the forward model. It was felt better to leave that judgement to the retrieval.

The figure presents the kernels for the retrievals with the greatest and least degrees of freedom (cases 3 and 6), showing both have sensitivity to δ , ϕ_{\parallel} , ϕ_L , Δ , ζ ,

⁶This corrects for the observed optical thickness and so will overestimate C near the surface.

⁷Application of this method requires simulating an elastic profile. For simplicity, a constant lidar ratio of 41 was assumed. Hence, these comparisons are favourable to the Wandinger method.

⁸(2.1.17) tends to slightly underestimate $\bar{\beta}$ such that (4.1.2) will slightly overestimate C . In the absence of noise, this produces the small positive bias Wandinger result (red).

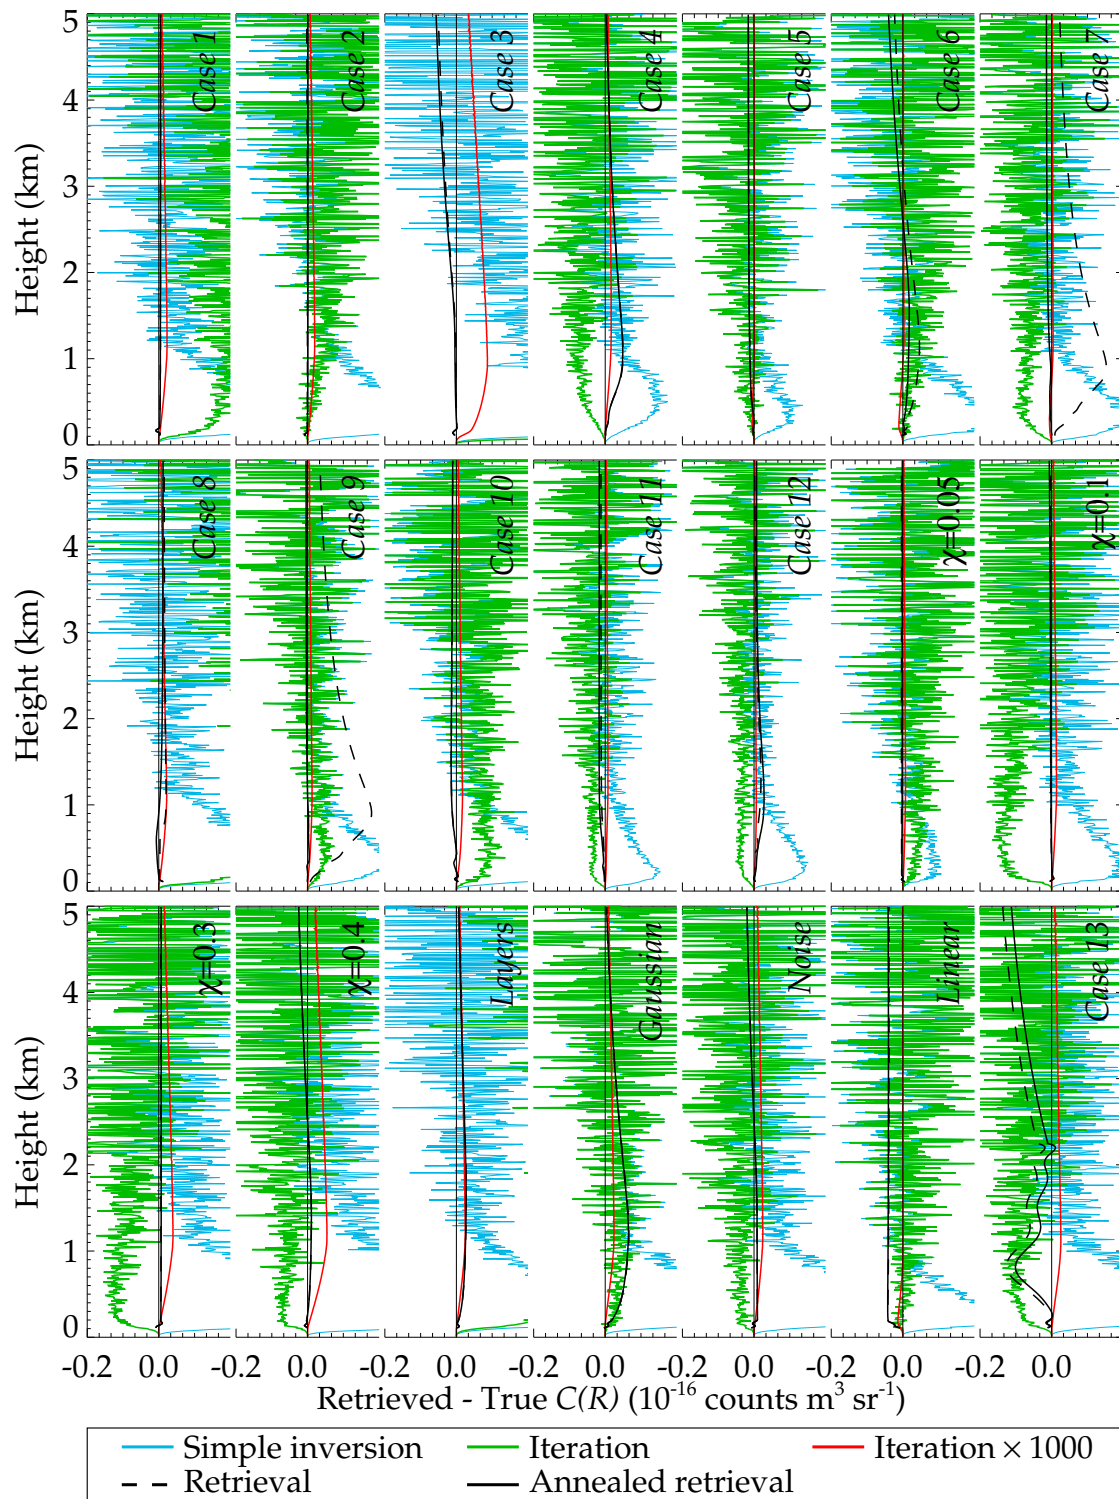


Figure 4.24: Comparison of the proposed retrieval to the technique of Wandinger and Ansmann [2002], showing deviation of retrieval from simulations for the data of figures 4.10, 4.15, and 4.17.

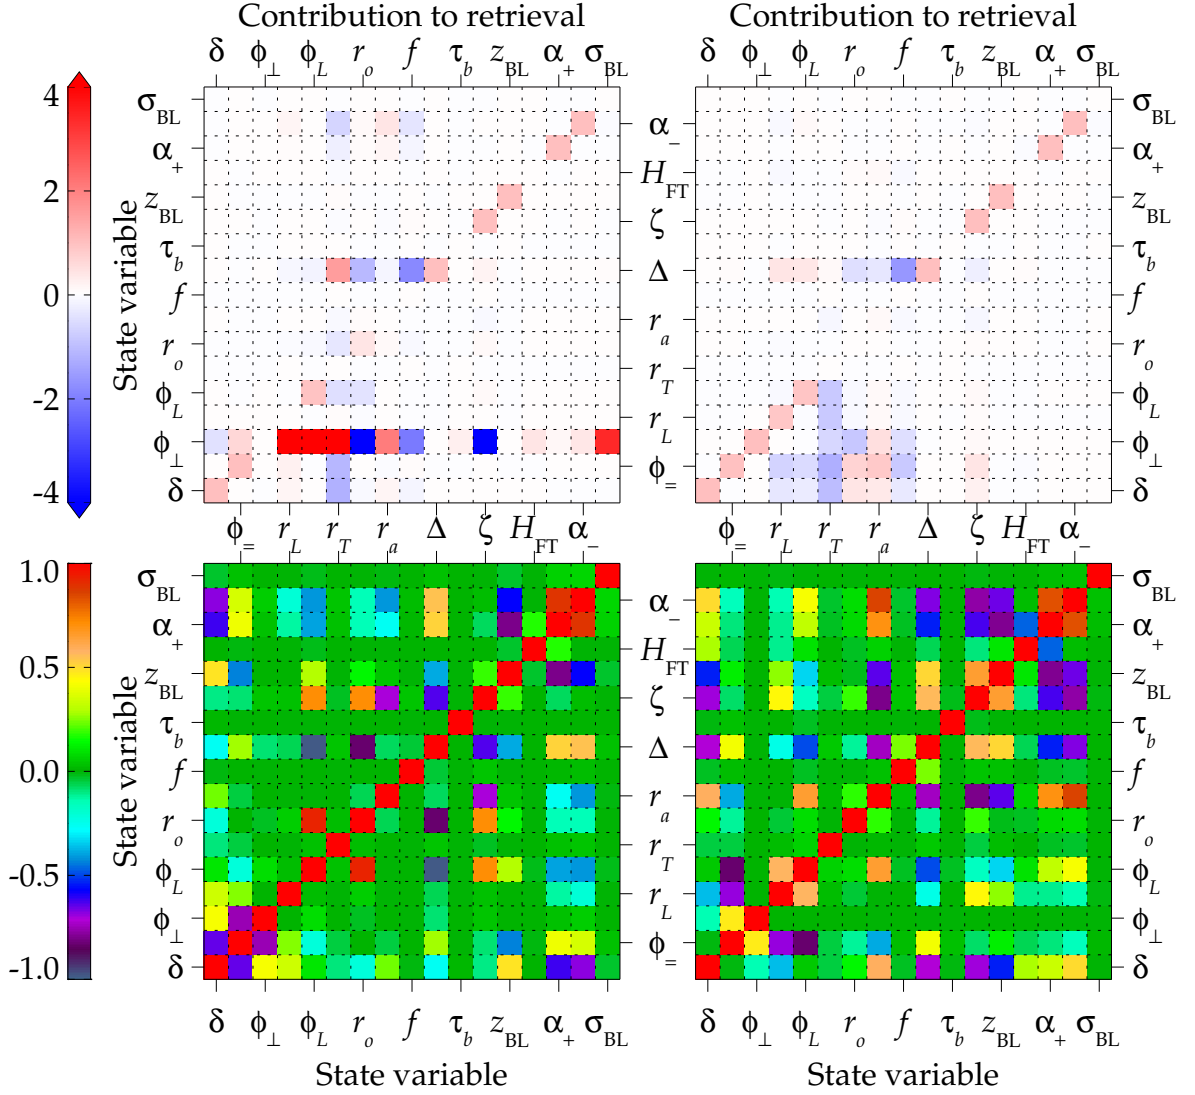


Figure 4.25: **Top:** Normalised averaging kernel, $(\mathbf{x}^{T\frac{1}{x}}) \mathbf{A}$, showing the contribution of each 'true' value of the state vector (x-axis) to the retrieved value (y). **Bottom:** S_x , normalised to give a correlation matrix. **Left:** Case 3, with the minimal degrees of freedom 8.4. **Right:** Case 6, with the maximal degrees of freedom 9.8.




Figure 4.26: Contributions to the total error variance from forward model parameters for case 7.

z_{BL} , and α_{\pm} . The more successful retrieval can also estimate ϕ_{\perp} and r_L . Hence, the published method was correct in its conclusion that ϕ_{\perp} was not the best choice of variable, with ϕ_L being a clear replacement. These elements are still under-constrained as some off-diagonal elements are non-zero, most significantly for ϕ_{\perp} and Δ . Interestingly, the retrieval with more available information uses the telescope dimensions to contribute to the retrieval of δ whilst the less well-constrained retrieval does not. This use of correlations between variables likely explains its greater degrees of freedom.

The covariance matrices are rather more difficult to interpret in this form. It is worth noting the high correlations between α_+ and α_- , which both drive the aerosol optical thickness, or between ϕ_L , r_o , and ζ , which all control the magnitude of the signal. Similarly, a strong anticorrelation exists between ϕ_L and Δ as they both affect the fraction of light coupled onto the detector, but in different senses. Case 6 (bottom right) bears the hallmarks of a difficult aerosol retrieval due to the large anticorrelation of α with z_{BL} , ζ , and Δ .




Figure 4.27: An illustration of the impact of laser jitter on the calibrations retrieved for cases 1–7, showing the profiles resulting from 10 random perturbations of $\phi_{||,\perp}$.

4.6.2 Parameter error

This retrieval is fairly computationally expensive, especially when annealing is applied⁹. As such, running with any inclusion of parameters errors was not deemed feasible (though they are included afterwards into the final error estimate, see Section A.4). Regardless, provided these errors are not too large, the solution does not appear to be affected as, unlike the retrieval of Chapter 2, τ_d is less important and so introduces little correlation between bins¹⁰. Figure 4.26 shows the contribution to S_ϵ from each parameter, assuming the parameter variances presented in Chapter 2. Somewhat surprisingly, R_0 is the most important parameter by an appreciable margin and would need to be known to 1 m for its contribution to be less than measurement noise. That is a fairly stringent limitation that, though theoretically achievable, would require laboratory measurements that are not attempted here. As with the previous retrieval, therefore, the uncertainty in the retrieved products is dominated by uncertainty in the parameters of the retrieval rather than native measurement error (a source of error not commonly discussed in the literature).

One final matter is laser jitter. Even in a laboratory environment, it is well

⁹A retrieval takes 3–5 min, which is doubled when evaluating parameter error.

¹⁰Errors of nanoseconds in τ_D and tens of metres in R_0 are sufficient to remove any information content from the measurement.

known that the alignment of a laser will drift over time due to slight shifts in the laser cavity and optical components from vibrations and thermal expansion. This drift is of order 0.1 mrad, which is similar to the uncertainty in $\phi_{\parallel,\perp}$, such that the error in any calibration will increase when applied to observations well before or after its own evaluation. Figure 4.27 shows the impact of adding random perturbations of this order to the calibrations retrieved from cases 1–7. The impact is small for the well-aligned cases, with the error only important below 300 m where there are already significant uncertainties. However, in cases such as 4 and 7, the effect is significant. As these drifts are entirely unavoidable, this highlights that it is both important to calibrate regularly¹¹ and to properly align the system to minimise the impact of drift.

4.7 Summary

An optimal estimation retrieval scheme for the calibration function of a Raman channel of a lidar system is proposed. This assumes the extinction profile conforms to an idealised form whereby it is constant through a well-mixed PBL, sharply decreases at its top, and decreases exponentially above that. The overlap function is modelled by the optical model outlined by Halldórsson and Langerholc [1978], optimised for numerical evaluation. *A priori* estimates are mostly derived from simple measurements of the system and models of the global aerosol distribution.

Various configurations of the forward model are possible for this retrieval, depending on the assumed form of the aerosol profile and the possibility of logarithmic retrieval of some state vector elements. A subset of these were chosen as they presented minimal RMSD between their retrieved C and that simulated in various circumstances. Different configurations were found to be appropriate for well and poorly aligned systems. However, it is also important that the retrieval demonstrate minimal sensitivity to the input parameters, such as Γ_0 and the retrieval's first guess. Only one retrieval was both accurate across a variety of circumstances and insensitive to its parameters — where aerosol extinction is modelled by an error function while retrieving all variables except ζ linearly.

The retrieval was shown to be reasonably insensitive to the aerosol profile (for $\chi_\infty < 0.3$), including where α did not conform to an idealised profile. Some difficulties remain in the retrieval of poorly-aligned calibrations or where the laser beam shape is Gaussian. These can mostly be avoided by annealing the solution

¹¹For example, ARM sites automatically calibrate every 14 hours [Goldsmith et al., 1998].

— whereby the first guess of δ , $\phi_{||}$, and Δ for the retrieval is randomised and the retrieval repeated several times, from which the physically-consistent solution with lowest cost is selected.

In these circumstances, the retrieval returns expected errors of 1–2 % without considering parameter error. Systematic forward model errors of a similar magnitude should be included to account for inconsistencies in the extinction profile, laser beam shape, and other optical effects. Considering errors on the model parameters gives a total uncertainty of 5–15 %, depending on the quality of the data. This is comparable to existing techniques, which have a quoted accuracy of $\sim 5\%$ in ideal circumstances but, unlike the method proposed, give no information about its variation with height nor the impact of any iteration scheme applied. The proposed technique has the added advantage that it is applied to measurements routinely made by any lidar — vertical observations of a well-mixed PBL — and so can be retroactively applied to any data set without having to perform additional calibration measurements. For systems where the calibration is already well-known, that information could be used as an *a priori* and first guess for this retrieval, which should then be able to reduce the uncertainty in the existing estimate and provide a smoother, more physically sensible result.

Chapter 5

Application to real data

It now remains to apply the retrievals developed in the previous chapters to real observations and compare their results to other methods of analysis. The degree to which the calibration techniques of Chapters 3 and 4 provide results useful for the extinction/backscatter retrieval will be of particular concern. This chapter will primarily concentrate on data collected with the CUV as those observations are better constrained and possess higher SNR than those from RACHEL.

5.1 Chilbolton ultraviolet lidar

For this research, two types of data were available from the CUV:

- Profiles of the attenuated backscatter coefficient (ABC) and relative humidity derived by CFARR after any period of observation and then published at the BADC; and,
- The raw data from 31 days of observations, with no additional processing.

As outlined in Section 1.4, the CUV is a more powerful system than RACHEL in many ways, but it is intended for daytime water vapour profiling such that the PMT discriminator thresholds and gains are optimised for analogue observations. Hence, the data will not always be ideal for aerosol sampling as desired here. Some of the technical details required for the proposed analyses are not known. With the calibrated ABC profiles this is less important, but raw data require careful consideration.




Figure 5.1: Raw elastic observations by the CUV shortly before noon on 2 Mar 2010, illustrating that for the PC mode the signal begins in bin 6 and 10 for the analogue mode. Other variations are mostly due to fluctuations in the laser energy.

5.1.1 Vertical offset

The ABC profiles are reported as 1,760 7.5 m bins starting at 200 m derived from analogue measurements only. Raw profiles are presented on a 10×1800 grid. No dedicated measurement of the vertical offset for these data could be found; it will have to be inferred from the data. Firstly, the 7.5 m bin length is an approximation as the ADC operates at 20 MHz. Using an accurate value of c therefore gives a bin length of 7.4948 m, which is admittedly a small error (about 10 m over the entire profile).

The vertical offset is estimated by considering the first bin to contain signal. Shown in figure 5.1, a very strong signal is registered in bins 6 and 10 for the photon counting (PC) and analogue modes, respectively. Taking these as the first bin (offsets of -37.5 and -67.5 m, both with an uncertainty of 8 m) gives profiles identical to the BADC product and good correlation with other lidar measurements (fig. B.10). The difference between the channels derives from the time necessary to operate the ADC and the bandpass of its pre-amplifier. Figure 5.1 also highlights a ringing in the first few dozen analogue bins. As in CFARR's analyses, these bins will be ignored.

5.1.2 Measurement uncertainty

Throughout the simulations presented in previous chapters, measurement noise was assumed to be Poisson distributed such that their variance could be approximated by the observation itself. This implicitly assumes that the observations correspond to a discrete count and that the dark current and thermal excitations are negligible. For the PC mode, this is reasonable as these sources of noise are efficiently filtered by the discriminator.

However, it does not hold for the analogue channel. Such observations are roughly Poisson distributed despite variations in pulse height and a DC offset [Fried, 1965], but their variance is not merely equal to their magnitude. The analogue voltage V should be a linear function of the number of photoelectrons U ,

$$V(R) = c_0 + c_1 U, \quad (5.1.1)$$

where $c_{0,1}$ are constants. As U is Poisson distributed, the measurement variance can be approximated by,

$$\sigma_V^2 = c_1^2 \sigma_U^2 = c_1^2 U = c_1(V - c_0). \quad (5.1.2)$$

This can be investigated by considering multiple measurements of a constant illumination, from which the mean and variance can be estimated. Such observations are available at the top of any lidar profile, where the solar background is the dominant signal in several bins. Figure 5.2 plots the variance against mean of the top 100 bins for all profiles from 2 Mar 2010 (sampling clear and cloudy conditions, day and night, for maximal dynamic range). At low signal levels, the PC measurements have equal mean and variance, as expected, but diverge at greater count rates as the detector's nonlinearity becomes important. It should be noted that uncorrected data overestimate the true variance.

The analogue Raman channel demonstrates a linear relation with high correlation, shown by a linear fit¹ of (5.1.2). The elastic channel captures little solar radiation and so samples an insufficient dynamic range to make a reasonable fit. As those measurements are consistent with the Raman observations, the Raman estimate will be used for both channels.

¹Neglecting errors, as it is not immediately obvious what they would be.

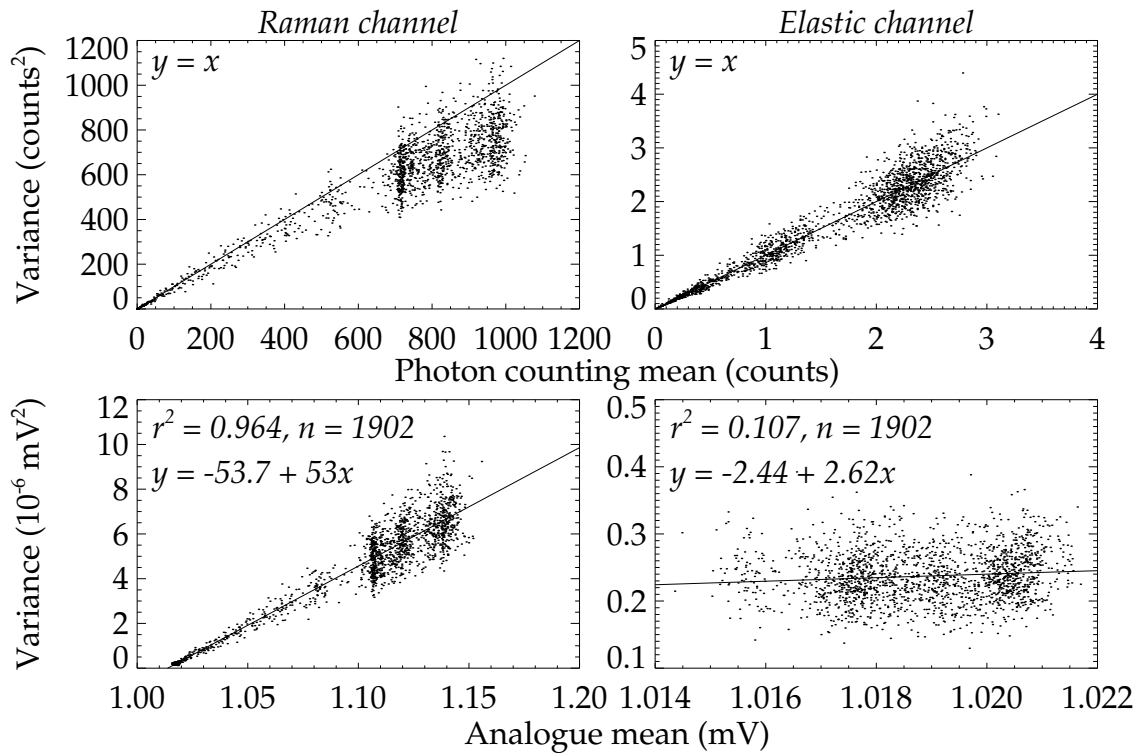


Figure 5.2: Variance versus mean of observations of the solar background on 2 Mar 2010. **Left:** Nitrogen Raman channel. **Right:** Elastic channel. **Top:** Photon counting observations over the one-to-one line. **Bottom:** Analogue observations over a linear fit to (5.1.2).

5.1.3 Calibration function

Five days of CUV observations have been selected from which the calibration function will be estimated. They will be considered in turn to highlight the strengths and weaknesses of the data and to compare the various calibration techniques. Figures B.12 and B.13 show the ABC during these days to emphasise the locations of clouds, variations in PBL height, and the regions over which data are averaged to produce a test profile for analysis.

In their own analyses, CFARR apply an overlap function determined from 30 minutes of analogue data on 21 Sep 2005. A calibration factor is then determined manually after each day's observations from comparison against a clean atmosphere. This factor is generally left constant, but has been changed a few times over the years after alterations to the apparatus. This fails to account for aerosol optical thickness, though days with higher χ_∞ are discounted from calibration. If the instrument becomes clearly misaligned, efforts will generally be made to realign it, which is assumed to return the instrument to its known overlap function




Figure 5.3: Comparison of C^{ra} retrieved for 21 Sep 2005, normalised between 4 and 5 km, showing the calibration with absolute lowest cost in orange and that with $\alpha_+ \geq \alpha_-$ and minimal cost in red (solid lines for the analogue data and dashed for PC). Simple inversions (5.1.3) of the analogue and PC profiles are shown in green and blue, revealing an inconsistency between them. The CFARR function is shown in black.

(with the misaligned data removed from the published products).

The data from which that correction was determined present an obvious starting point. That day was cloud-free, though a single profile at 11:49 contains a small maximum at the top of the PBL and is omitted. There was also a step-like increase in laser power at 11:36; data before then will not be considered. Variations in the ABC clearly reveal mixing of aerosol in the PBL, but these are less than 10 % so the scene is considered homogeneous. These observations pre-date the AERONET station at the site so there are no independent measurements of χ_∞ . The *a priori* value of 0.14(28) will be assumed.

Figure 5.3 presents various overlap functions derived from these observations. It shows a simple arithmetic inversion of the test profile,

$$C(R) \simeq \left[\frac{\varphi(R)}{1 - \tau_d \varphi(R)} - E_B \right] \frac{R^2 \exp[(1 + \lambda^{\text{el}}/\lambda^{\text{ra}})\chi_\infty + \chi_m]}{MN(R)}. \quad (5.1.3)$$

For these electronics, the analogue and PC data give different curves (green and blue). This would usually be attributed to detector nonlinearity, but the analogue mode presents a *smaller* peak value than the PC. Conventional wisdom is that, at the magnitudes observed here, the analogue mode should give a linear measure of the photon count rate and the PC mode should be attenuated by dead time effects; that is clearly not true here. The discrepancy has been attributed to the optimisation of the system for analogue water vapour sounding, such that the PC mode is

not behaving entirely as expected, but no precedent for that has been found in the literature. In the absence of additional information, the data will be processed ‘as is’ with correction only to the PC measurements.

Secondly, there is a slight difference between CFARR’s overlap function and the simple inversion. This is because CFARR used a sliding-window fit on the curve². Regardless, the retrieval from analogue data agrees fairly well with CFARR’s product. The retrieval was annealed over 100 first guesses³ of which two results are presented — the solution with the absolute lowest cost and the solution with lowest cost that used a physically reasonable aerosol profile (where $\alpha_+ \geq \alpha_-$). The latter is qualitatively very similar to the existing estimate but, in detail, diverges by a few percent with height and by up to 50 % below 500 m where the aerosol is most important. PC observations present an overlap function that gently peaks at 1.6 km and is always greater than the analogue calibration. In either case, the selected solutions is consistent with others in the ensemble provided they retrieve similar χ_∞ .

As there are no measurements in addition to these profiles, the only other applicable method is that of Wandinger and Ansmann [2002], being (4.1.2). Shown on the left of figure 5.4, there is a reasonable agreement between that method and the retrieval for the PC observations (blue against dashed lines), but not the analogue. This is misleading for several reasons. The retrievals are plotted normalised by their mean between 4 and 5 km. The Wandinger technique is bounded by assuming that the same region is free of aerosol. Though this appears to be a reasonable assumption⁴, the magnitude of the overlap function is highly dependent on the range chosen due to the substantial noise there. Hence, the agreement shown isn’t necessarily meaningful as there are two components to a calibration — its shape and its magnitude. The right-hand plot illustrates this by normalising all the curves with their mean between 1 and 1.2 km. This demonstrates that the various techniques are giving similar functional forms beneath 1 km, but tend differently above that.

Secondly, the shape of the Wandinger solution is weakly dependent on the lidar ratio assumed. This is determined by comparing the χ_∞ derived from the elastic

²The correction was confirmed by comparison to the 353.9 nm channel, which suffers virtually identical extinction, but no backscatter. Care must be taken due to the temperature dependence of that channel.

³In order to decrease the number of clearly incorrect solutions produced, the ensemble was constructed in sets of ten. For each, the C resulting from 100 x_0 were evaluated and the ten with the lowest initial cost would be used to seed a full retrieval.

⁴In that the Fernald-Klett solution in this region is not affected by the chosen lidar ratio.




Figure 5.4: **Left:** As fig. 5.3, but showing the result of the Wandering calibration method (4.1.2) rather than the simple inversion. **Right:** The same data, but instead normalised by their mean between 1 and 1.2 km.

channel after overlap correction to its known value. In this case, the analogue and PC observations select values of 185 and 85.6 sr, respectively. The former is clearly wrong and using a more realistic value (such as 60 sr [Müller et al., 2007]) brings the retrieved analogue solutions into closer agreement with Wandering. However, it would be preferable to consider a better constrained scene.

The 18th of July, 2007 presents similar conditions to the previous scene but with AERONET measurements (averaging 0.119(12) from level 1.0). Shown in figure 5.5, the analogue and PC modes continue to disagree, but both indicate a smaller calibration than the previous profile, especially near the surface. Whereas the PC profile previously appeared to reach full overlap around 1.2 km, it now varies in a manner more similar to what CFARR assume. The analogue profile never reaches full overlap. The Wandering solutions are consistent with the retrievals, but that is mostly because those are very uncertain due to noise above 4 km.

Further examples were drawn from 10 Aug 2007, which presents an hour of clear-sky observations with level 2.0 AERONET measurements throughout (0.36(1) average). Though the optical thickness is large, the ABC profile remains smooth, indicating the PBL is well-mixed. Three test profiles were produced — one averaging the entire period and one each from the first and last 23 minutes. The overlap

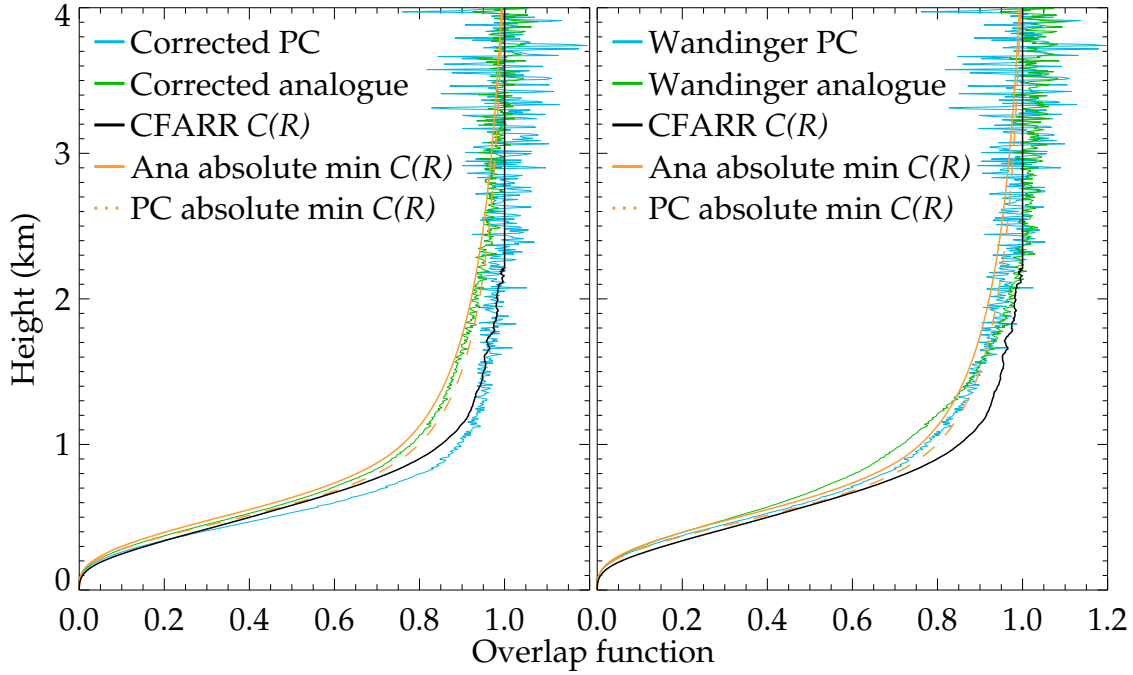


Figure 5.5: As figures 5.3 (left) and 5.4 (right), but for 18 Jul 2007. Each is normalised over 4–5 km. The absolute minimal solution is also physically consistent.

functions derived from all five test profiles are compared in figure 5.6. Plotting the solutions normalised at height appears to give several functions, with the retrieval selecting only two of those. However, by normalising around 1 km it is clear that the solutions are in fact quite similar. Though 18 Jul 2007 gives larger C at all heights, the retrieval has captured that difference.

The retrieved functions do not reach and maintain a maximal value above some height, as would be preferred. The spread of Wandering solutions illustrate how little information is available at the top of the profile to constrain the solution. Hence, the retrieval is not penalised for using solutions with excessively small ϕ_L which causes decreasing C as R tends towards infinity⁵. Future work may wish to more tightly constrain the calibration function model to give more traditional behaviour as the information content decreases.

The analogue results are less appealing, but this can again be attributed to unrealistic values of the lidar ratio — 185, 135, 95.8, 119, and 84.9 sr. Once again selecting more reasonable values gives more a favourable comparison, but that remains an unsatisfying solution.

⁵This corresponds to cases 5 and 9 of fig. 4.5 where a small laser beam falls within a large region of incomplete focus onto the detectors.




Figure 5.6: Comparison of retrieved and Wandinger estimates of the overlap function from test profiles in 2005–7. CFARR's product plotted in black for reference. **Left:** PC measurements. **Right:** Analogue measurements. **Top:** Normalisation between 4 and 5 km. **Bottom** Normalisation between 1 and 1.2 km.

The discrepancy between the PC and analogue measurements ceased after the detectors and electronics were replaced in Spring 2009. However, as shown in the right-most plots of fig. B.13, the applied nonlinear correction is no longer entirely sufficient. This is because τ_d was taken as lying somewhere between the minima of fig. 3.13 rather at the optimal value for these data. Parameter error should account for this discrepancy. In addition, even after correcting for the measured laser power, discontinuities remain in the ABC over time (fig. B.11a). On the one hand,




Figure 5.7: As fig. 5.6, but for 2 Mar 2010.

the laser energy measurement appears to be accurate as the signal in the lowest bins is consistent (see Section 3.2.2). On the other, the ratio of the elastic to Raman observations does not show such discontinuities (fig. B.11b). This may be due to a poor coupling of the beam onto the power meter, giving only an approximate measure.

On Tuesday, March 2nd 2010 nothing particularly notable was observed, though there are unusually measurements at night⁶. Cloud develops above 4 km over the course of the measurement period and the PBL height varies from a minimum of 800 m at 12:00 to a maximum of 1.4 km after 20:00. Though χ_∞ is observed to

⁶The CUV must be attended during operation.




Figure 5.8: As fig. 5.6, but for 16 Mar 2010.

remain 0.17-18, there are obvious inhomogeneities in the ABC.

Regardless, four test profiles averaged over 80 minutes were produced (a sufficient duration for the Wandinger method to give consistent answers). These are increasingly contaminated by cloud above the range considered by the retrieval, such that the fourth profile has a less accurate estimate of χ_∞ due to a lack of AERONET measurements. Shown in figure 5.7, it is immediately obvious that the calculated overlap function is different to that expected; it no longer achieves ‘full’ overlap. The Wandinger estimates from the first three test profiles are consistent with each other, whilst the fourth profile departs above 1 km. This may indicate a drift in the calibration but is more likely a failure of the retrieval in the presence of

cloud. As in previous cases, the retrieval is in good agreement with the shape of the Wandinger profiles for the PC measurements but is not capturing their magnitude. It fails on both counts with analogue observations. The analogue signal may still be nonlinear, merely less so than the PC, but the difficulties could derive from the lower information content available at the top of an analogue profile due to its greater background, E_B . The retrieval is then only sensible in those regions that possess enough information to constrain it.

The final day to be considered is 16 Mar 2010. Due to extensive, if thin, upper cloud, there are only three level 2.0 AERONET measurements available covering a short period around noon. Three profiles were selected — an hour of observations of a fairly stable and well-mixed boundary layer with several cloud-contaminated χ_∞ measurements around 10:00, a second hour of observations around noon encapsulating the level 2.0 measurements, and a twenty minute period within that hour that was removed from CFARR's published product as the alignment was perceived to have drifted. The results (fig. 5.8) clearly confirm that suspicion⁷. The retrieval is in good agreement with the iterative estimate for all PC data (fig. 5.9). The analogue retrievals remain problematic.

It is important to note that the elastic and Raman channels are expected to have different overlaps, under which conditions it isn't entirely clear what the Wandinger algorithm is estimating. Simulations appear to be dominated by the elastic overlap, but this is not universal. For the observations from 2010, this can be further investigated as a 355 nm EZ lidar was operational, against which the elastic channel can be calibrated. Figure 5.10 compares the answers resulting from each test case and reveals a number of interesting features:

- The calibrations retrieved from analogue and PC data are practically identical in cases 6, 8, and 10 (and are consistent in all cases). This is very reassuring as the two products must, in reality, be identical and any differences will result from improper correction of the data or noise.
- For analogue data, where the retrieval was not consistent with independent measures, the Wandinger solution closely resembles the estimate of the elastic overlap function⁸ (blue and green lines).

⁷Considering the retrieved values in tables B.5 and B.6, ϕ_\perp increases from zero to 1.13 μrad over the three profiles whilst ϕ_\parallel remains constant.

⁸Saying these curves are consistent isn't very meaningful, as the errors on the calibration against the EZ lidar have large and uncertain errors.

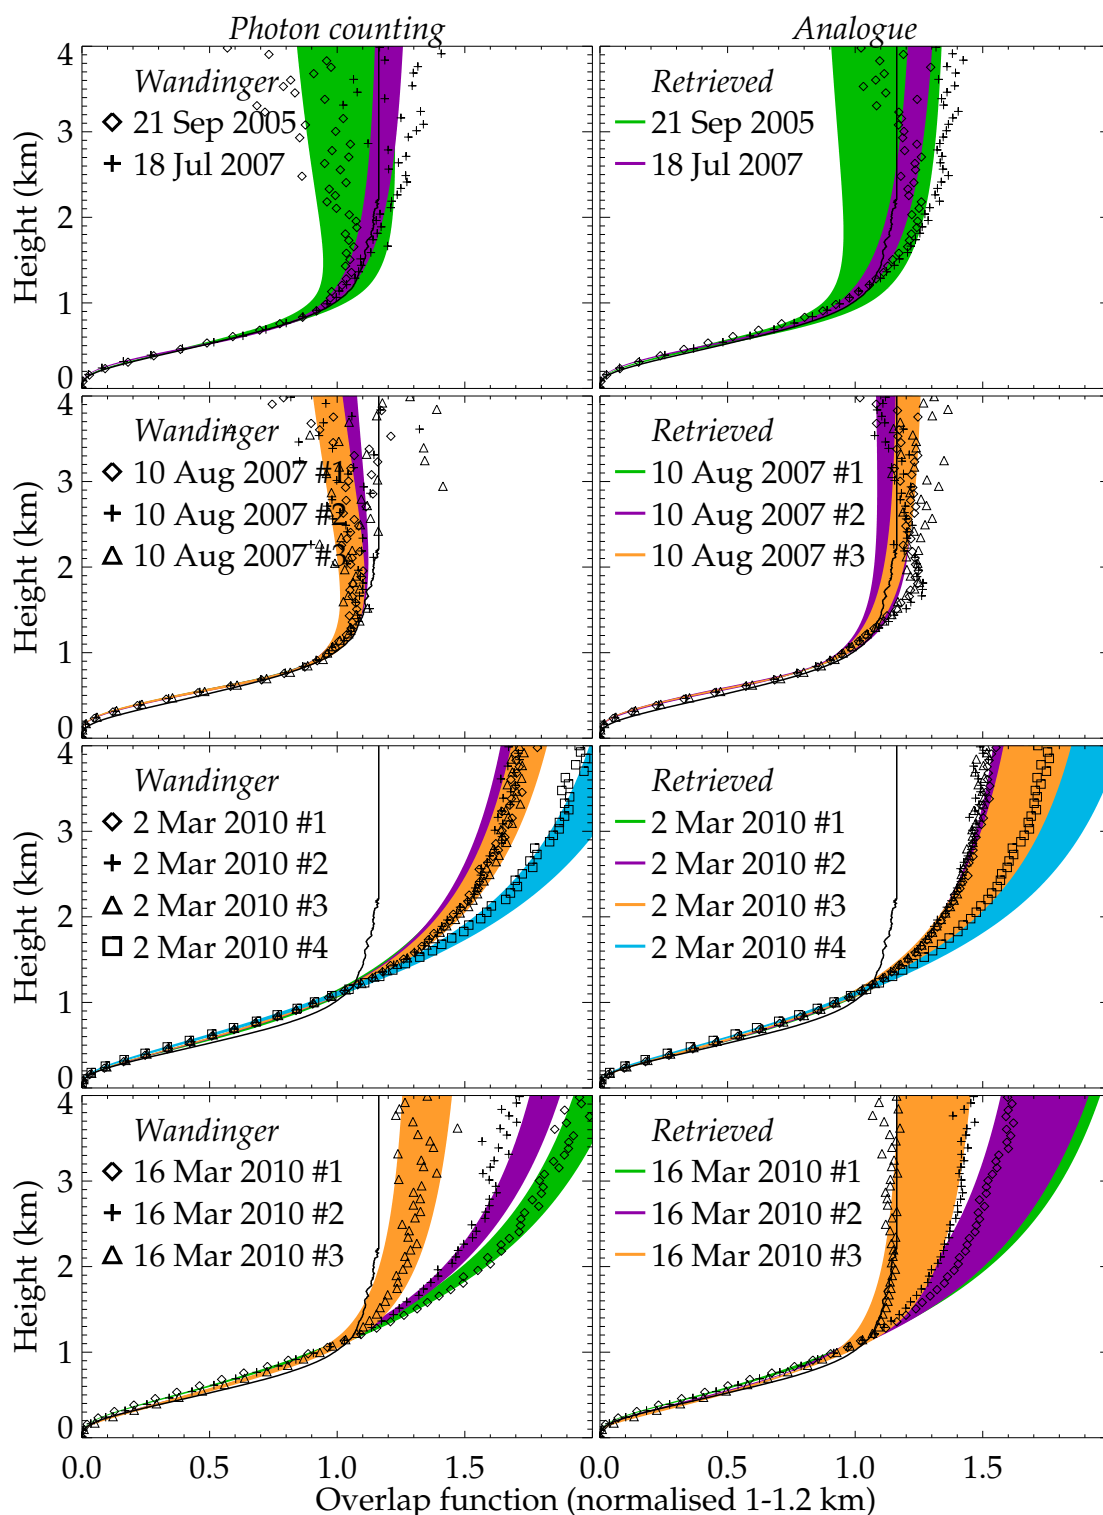


Figure 5.9: Illustration of calibration function retrieval errors compared against the estimates of the Wandering algorithm for the PC (left) and analogue (right) observations over the 11 test profiles.




Figure 5.10: Comparison of overlap functions derived by the Wandinger method, optimal estimation, and direct calibration against an EZ lidar for both the PC and analogue modes. The CFARR product is plotted as a dashed line for comparison.




Figure 5.11: Aerosol optical thickness derived from an arithmetic inversion of seven analogue profiles on 2 Mar 2010, using the retrieved calibration function (bottom) and that derived from the Wandinger method normalised over 4–5 km (top). The dotted line denotes 100 m.

- For PC data, the Wandinger solution closely agrees with that retrieved whilst the estimate of the elastic overlap is significantly smaller.

These taken together imply that the elastic and Raman channels do have slightly different overlap functions and, in that case, the Wandinger algorithm tends towards one of them (though it is not obvious which *a priori*). This by no means provides a conclusive validation of the method introduced in Chapter 4, but is the best possible with the data available and is hopefully sufficient to show that the proposed technique is at least as good as the existing methodology.

These calibrations are applied to other data in figure 5.11. Twenty second observations from 2 Mar 2010 were corrected with C retrieved from cases 6–9. The results aren't ideal — the earliest profile gives $\chi = 0.1$ at the surface whilst the latest gives -0.1 . These could be altered by adding a laser energy correction ne-

cessary to produce $\chi = 0$ at the surface, but that will break the correlation between χ_∞ and that observed by AERONET. However, such energy corrections also qualitatively improve the result from the elastic channel. It is left to the retrieval to deem if this is a reasonable solution. At the very least, though, these profiles now tend monotonically from a minimum to a maximum, which is not the case if the CFARR correction is applied. The scattering retrieval should, therefore, converge.

5.1.4 Extinction and backscatter

It is now possible to consider the retrieval of aerosol extinction and backscatter. For the remainder of this section, C retrieved from PC data will be applied to all channels, neglecting any differences between them as there is insufficient information to do otherwise. When applied to analogue observations, it will be scaled by the ratio of the retrieved calibrations $C^{\text{analogue}}/C^{\text{PC}}$ averaged over the top 300 bins. The scale factor for the elastic channels is determined by normalising against a clean region of the atmosphere, but the test data are too noisy to produce a consistent solution. Instead, five hours of observations are averaged to form the ratio,

$$\frac{1}{B\sigma_R} \frac{E^{\text{el}}(R) - E_B^{\text{el}}}{E^{\text{ra}}(R) - E_B^{\text{ra}}} \exp \left[\left(1 - \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \chi_\infty \right], \quad (5.1.4)$$

which is then averaged over 4–5 km. There is a greater uncertainty in χ_∞ by considering such an extended integration period, but the result is more independent of the height range selected. A more rigorous derivation of these quantities is an obvious area for future improvement. The additional error is around 4% as estimated by Monte Carlo techniques.

Figure 5.12 compares the profiles retrieved⁹ from the previous section's test data to those of the Fernald-Klett (2.1.10) and Ansmann methods¹⁰ (2.1.17) when the calibration function from test profile 6 is applied. In the PBL, the retrieved backscatter is very similar to that given by the Ansmann ratio. As the SNR decreases, the retrieval tends towards the Fernald-Klett solution. This is a proper response for the retrieval, giving answers similar to existing methods but tending from a two to one-channel retrieval as the available information decreases. This is also expressed in the averaging kernels which widen from 30 to 100 m (fig. 5.14).

⁹Using 150 bins between 0.1–5 km while retrieving $\ln \beta$ and B with $\Gamma_0 = 10^5$.

¹⁰Normalising between 4–5 km with the derivative evaluated over 150 m.

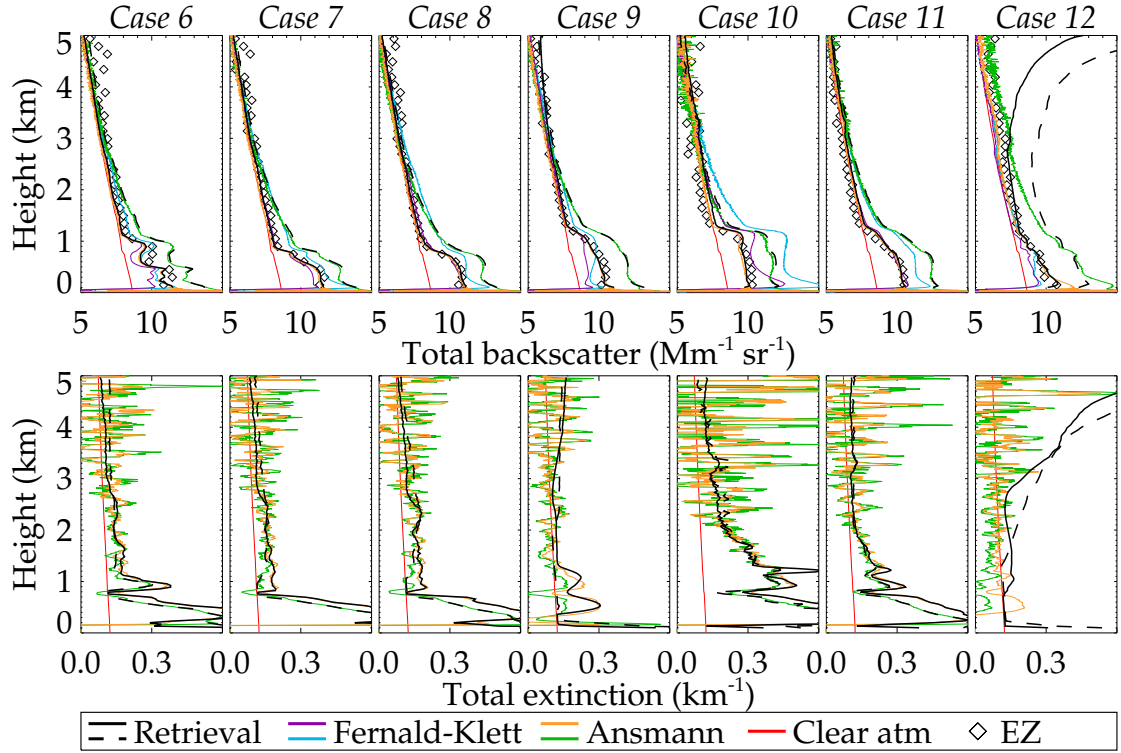


Figure 5.12: Various estimates of total backscatter (top) and two-way extinction (bottom) from the test profiles of Section 5.1.3 using C retrieved from profile 6. Retrieval from PC data (log mode) is plotted as a dashed line, with the result of the Fernald-Klett method in blue and the Ansmann method in green. For analogue data, these are solid, purple, and orange, respectively. The ABC reported by the coincident EZ lidar is shown in diamonds for comparison.

The retrieved extinction follows the Ansmann solution, though due to the influence of the *a priori* it gives a much smoother solution. The averaging kernels confirm that there is little information available in the free troposphere, but also show that the resolution in the PBL is better than observed with simulations: 100 m. A tendency to find $\alpha = 0$ at the top of the PBL is due to the number density profile. The radiosonde that morning recorded a step decrease in pressure at the top of the PBL, but as that is a low-resolution measurement, linear interpolation overestimates N there. A standard atmosphere does no better.

Retrieving in the linear configuration gives similar results, though that mode is more likely to give unbounded answers as the information content decreases (as in case 12 above). The inclusion of parameters errors reduces the information content, producing unbounded solutions in the free troposphere but doesn't greatly change values within the PBL (fig. 5.13). Therefore, parameter errors will only be considered after retrieval. These, as shown in figure 5.15, mostly behave as ob-

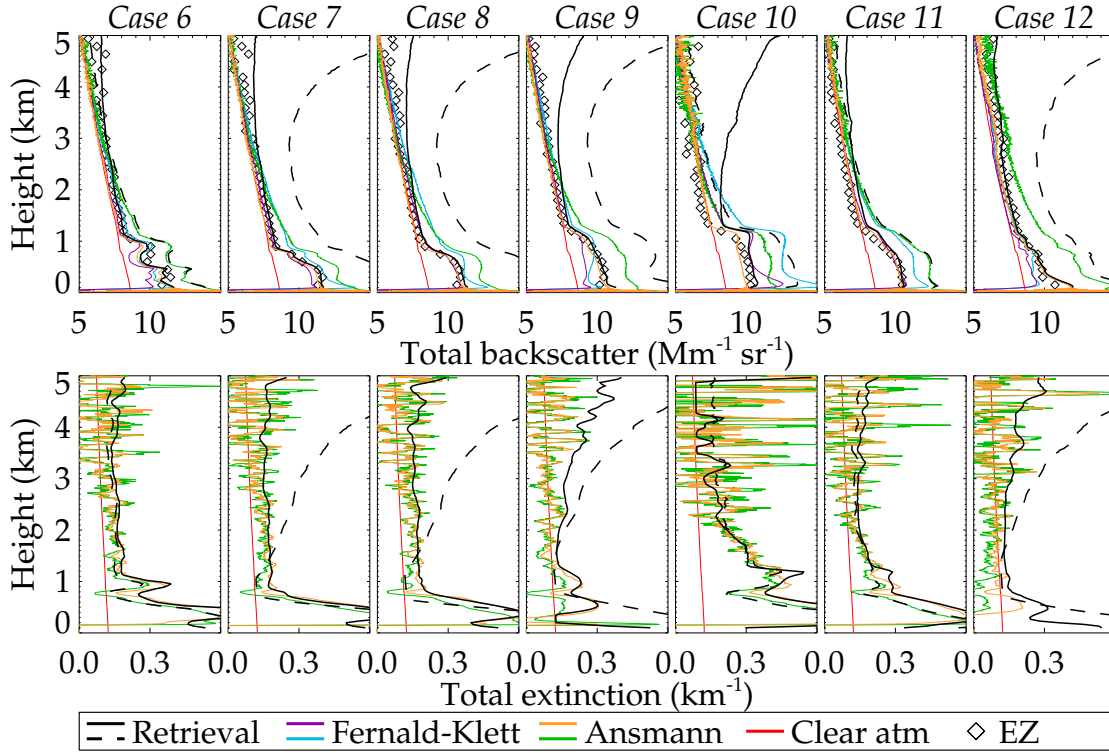


Figure 5.13: As fig. 5.12, but considering parameter errors during the retrieval.

served in Chapter 4 but with two exceptions:

- Larger background levels (being daytime observations) produce large correlations at the top of a profile¹¹. The two modes respond differently to this, with the logarithmic configuration reverting to the *a priori* covariance at the top of the PBL (as observed in simulations) and the linear configuration tending towards complete correlation (representing a more systematic error).
- The intercorrelation of β and α has changed (i.e. the quadrants of S_x that do not contain the diagonal). Bins above a point are still negatively correlated, but those below are only positively correlated in the linear mode and to a greater extent than previously. The exact reason for this is not clear.

Returning to figure 5.12, a large difference between the analogue and PC analyses is evident. This could indicate the nonlinear correction of the PC data is incorrect as the two modes tend towards similar answers and the discrepancy is evident in the Ansmann ratio, which is not sensitive to C . The ABC observed by

¹¹These are not due to parameter errors in E_B but express the retrieval's use of information.

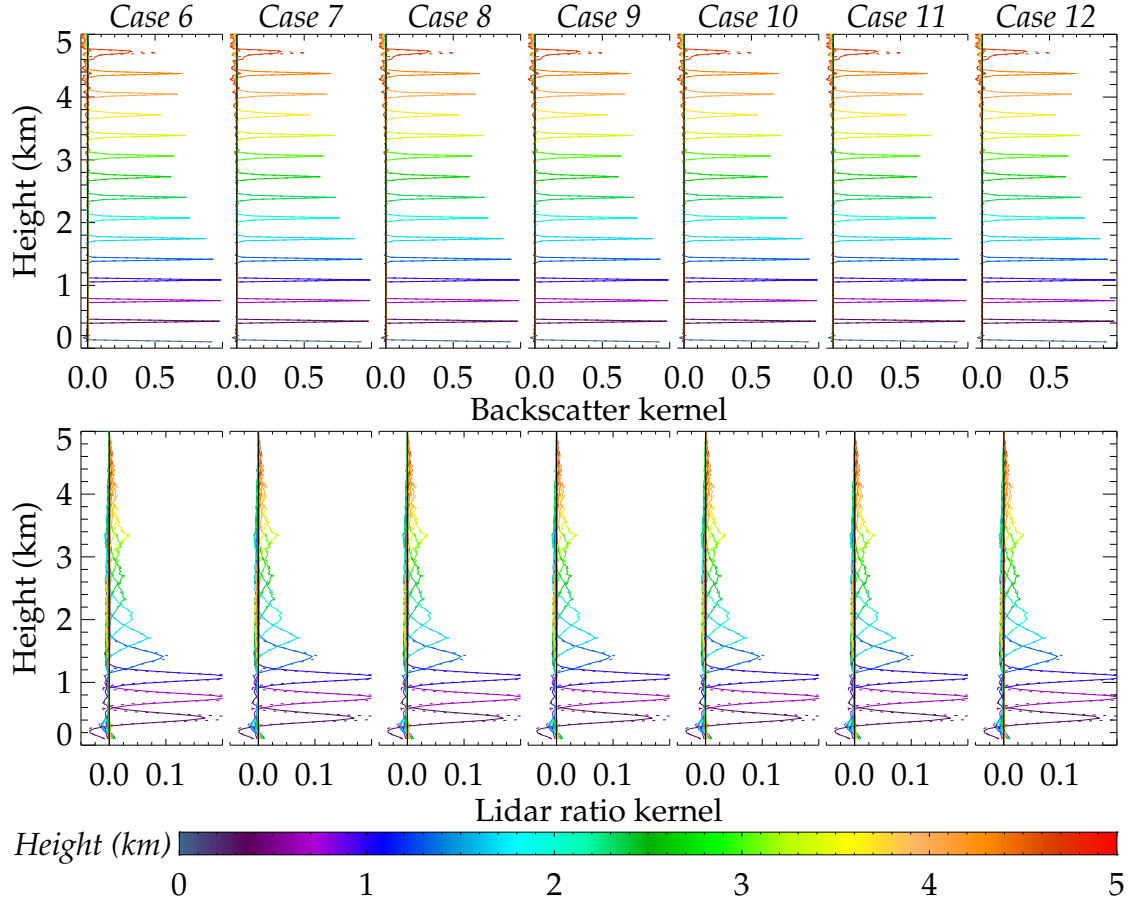


Figure 5.14: Averaging kernels for the retrievals of fig. 5.12, revealing resolutions of 30–100 m for backscatter and 0.1–1 km for the lidar ratio. **Solid:** PC measurements. **Dotted:** Analogue measurements, which are virtually identical.

the EZ lidar is shown for comparison, indicating the analogue analysis is nearer to the truth¹².

Figure 5.16 shows how the retrieval varies when different calibration functions are applied. Calibrations that were in good agreement with the Wandinger method (see fig. 5.10) give consistent results. There is some variation in their magnitude, but it is within the predicted errors; see fig. B.14. The lidar ratio is not constrained in the free troposphere, wandering to unrealistically large values. This is worse in the linear retrieval mode (fig. B.15), but as neither mode claims to be producing a successful retrieval there, the difference is moot. Figure 5.16 also reveals that though the analogue mode is closer in magnitude to the EZ lidar observations, it does not capture their shape as accurately as the PC retrieval. Further, its lidar

¹²The degree of agreement between the two will never be perfect as the EZ lidar product is attenuated by atmospheric transmission.




Figure 5.15: Errors for the retrieval of test case 10 by the linear and log modes (top and bottom, respectively). **(a)** Backscatter error. **(b)** Extinction error. **(c)** Backscatter autocorrelation. **(d)** Extinction autocorrelation. **(e)** β/α intercorrelation. **(f)** Log backscatter error. **(g)** Lidar ratio error. **(h)** Log backscatter autocorrelation. **(i)** Lidar ratio autocorrelation. **(j)** $\ln \beta/B$ intercorrelation.

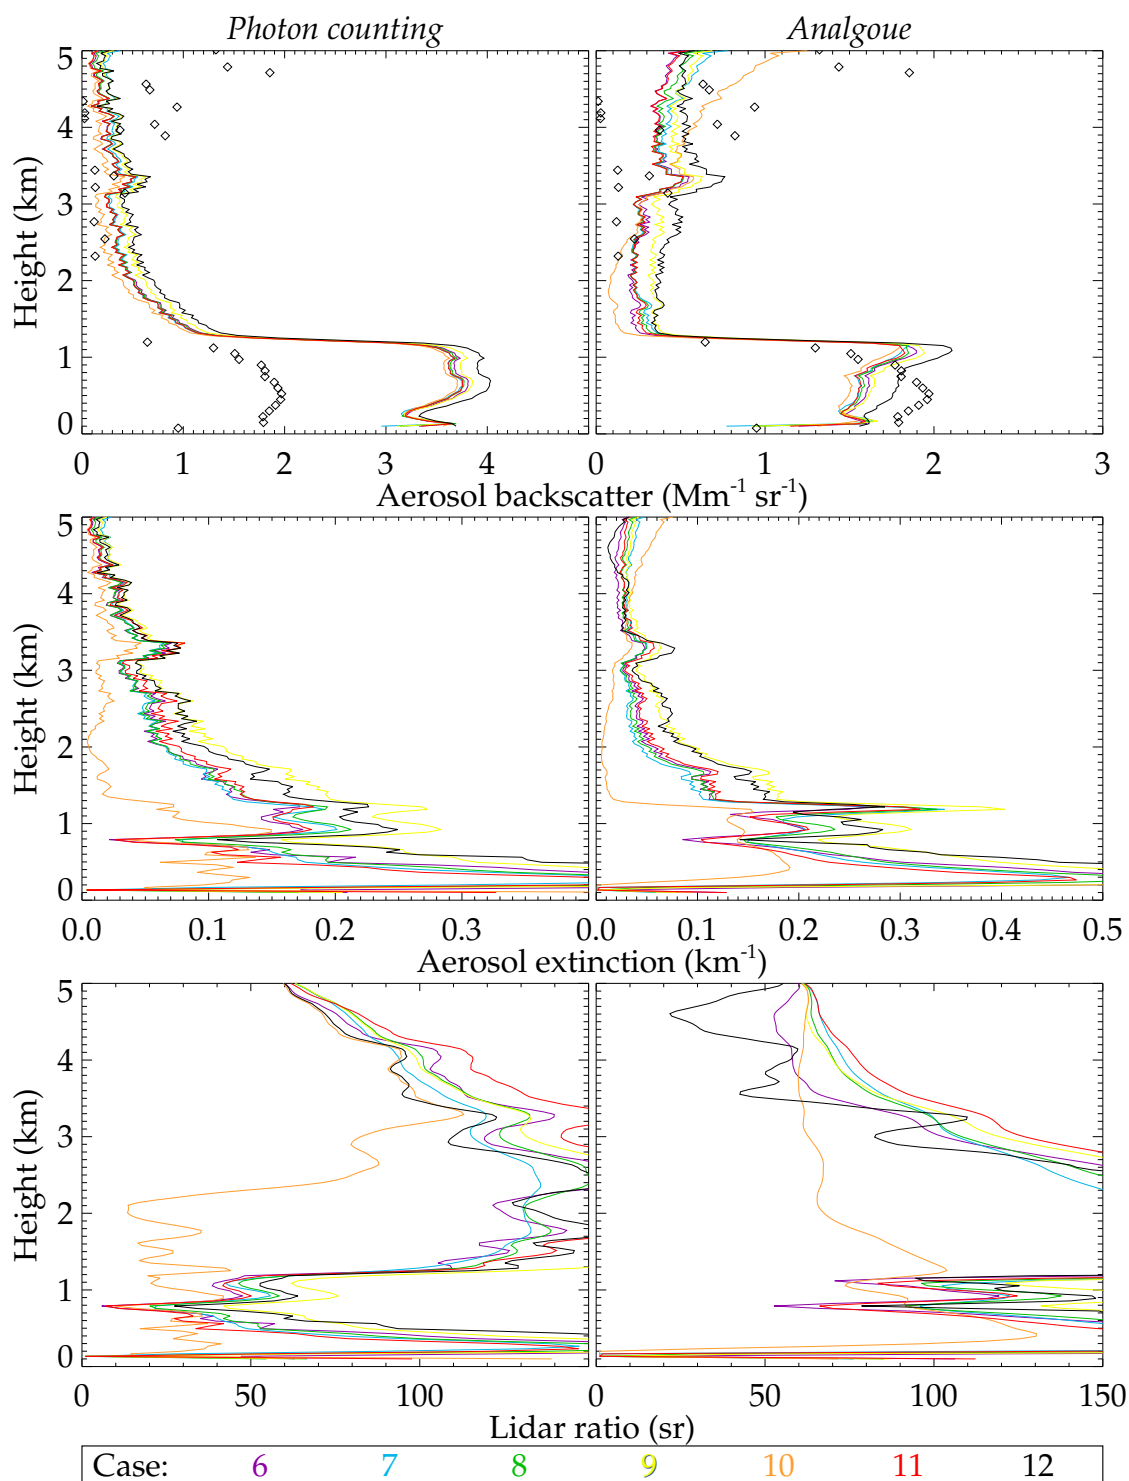


Figure 5.16: Comparison of the aerosol backscatter (top), extinction (centre), and lidar ratio (bottom) retrieved from test profile 10 when using various retrieved calibration functions (colour). The ABC reported by the EZ lidar is shown in diamonds for comparison.

ratios are unrealistically large, implying both the presence of nonlinearity in the analogue data and an overestimation of C^{ra} .

The calibration function retrieval has been shown to successfully estimate the shape of C and finds a magnitude consistent with existing techniques. As none of those is wholly satisfactory, further work is clearly required to better constrain the magnitude of C . The degree to which these issues are understood within the lidar literature varies. Goldsmith et al. [1998] developed an extensive methodology to estimate and validate all elements of the calibration of automated lidar systems whilst an equivalent paper by Althausen et al. [2000] merely states that their calibration is made against a clear sky and “the resulting error is less than 5 % to 10 %.” Though that range is probably valid, the authors fail to note that the error is dependent on the height range over which the reference is taken and is systematic (in the sense that every point in the profile will suffer an identical error).

Such failures, however, do not impact the ability to validate the extinction and backscatter retrieval against existing techniques. Figure 5.12 clearly shows that they are in agreement and that, as the SNR decreases, the retrieval produces a smoother answer with an estimate of its resolution. Validation in the strictest sense would compare against completely independent measurements but is not currently reasonable, especially considering the known difficulties in the intercomparison of lidar to other systems [Naud et al., 2003; Wiegner et al., 2006; Pahlow et al., 2006; Bradley et al., 2010]. It is acknowledged that the results to follow are strictly wrong, in that the calibration is known to be flawed¹³ and any conclusions drawn from them will not be scientifically defensible. However, they are the best that can be done with the resources available.

The retrieval is now applied to all observations from 2 Mar 2010. Various integration times were evaluated and found to give consistent results. Hence, only the highest resolution (1 min) is presented¹⁴. The backscatter field retrieved in the linear mode using PC data is shown in row (a) of figure 5.17. In many ways, the retrieval is performing excellently — the laser energy variations (highlighted in fig. B.11) have been mostly removed as the retrieved field is smooth, variations on the scale of tens of metres seen in the Ansmann ratio are also captured by the retrieval, and uncertainties in the PBL are 5–10 % (comparable to existing techniques).

¹³Further, the error due to the flawed calibration is not reflected in any products as when fairly estimated it eliminates any information content in the data.

¹⁴Interestingly, the error returned by the retrieval does not change with the integration time as parameter errors are dominant.

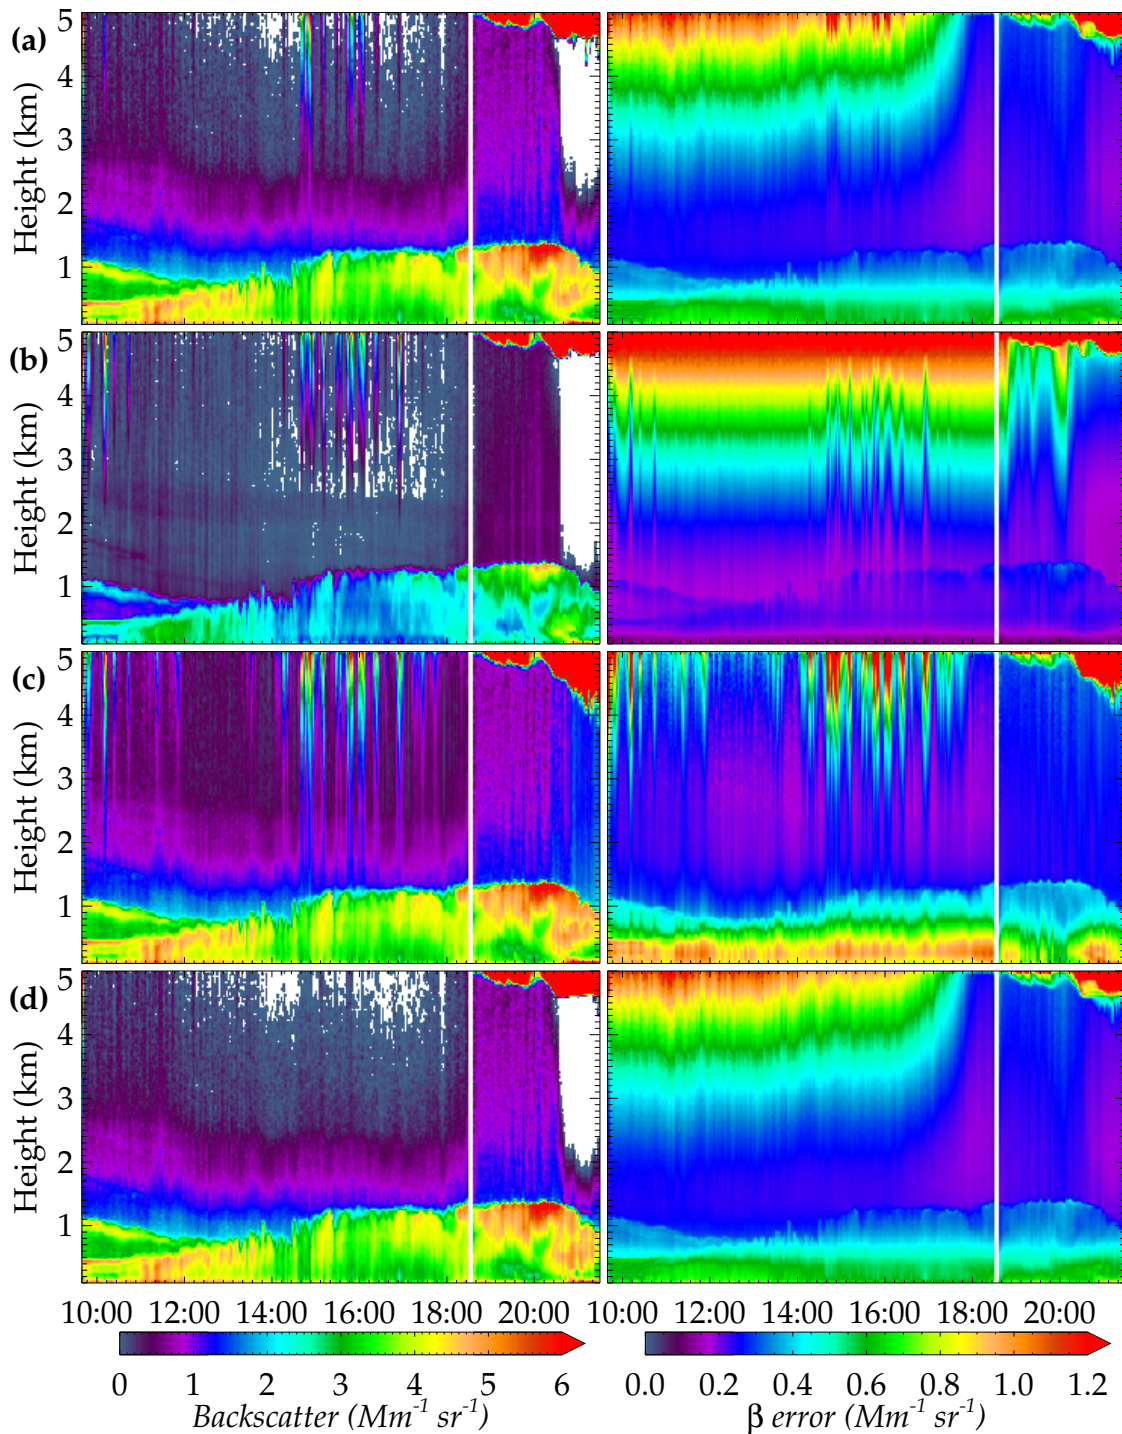


Figure 5.17: Retrieved aerosol backscatter (left) and its error (right) from 2 Mar 2010 at 33 m and 1 min resolution using the calibration derived from test profile 6. White indicates missing data or $\beta = 0$ (the retrieval's lower bound). (a) PC data, retrieving β and α . (b) Analogue data, retrieving β and α . (c) PC data, retrieving $\ln \beta$ and B . (d) As (a), but using each retrieval as the first guess for the next.

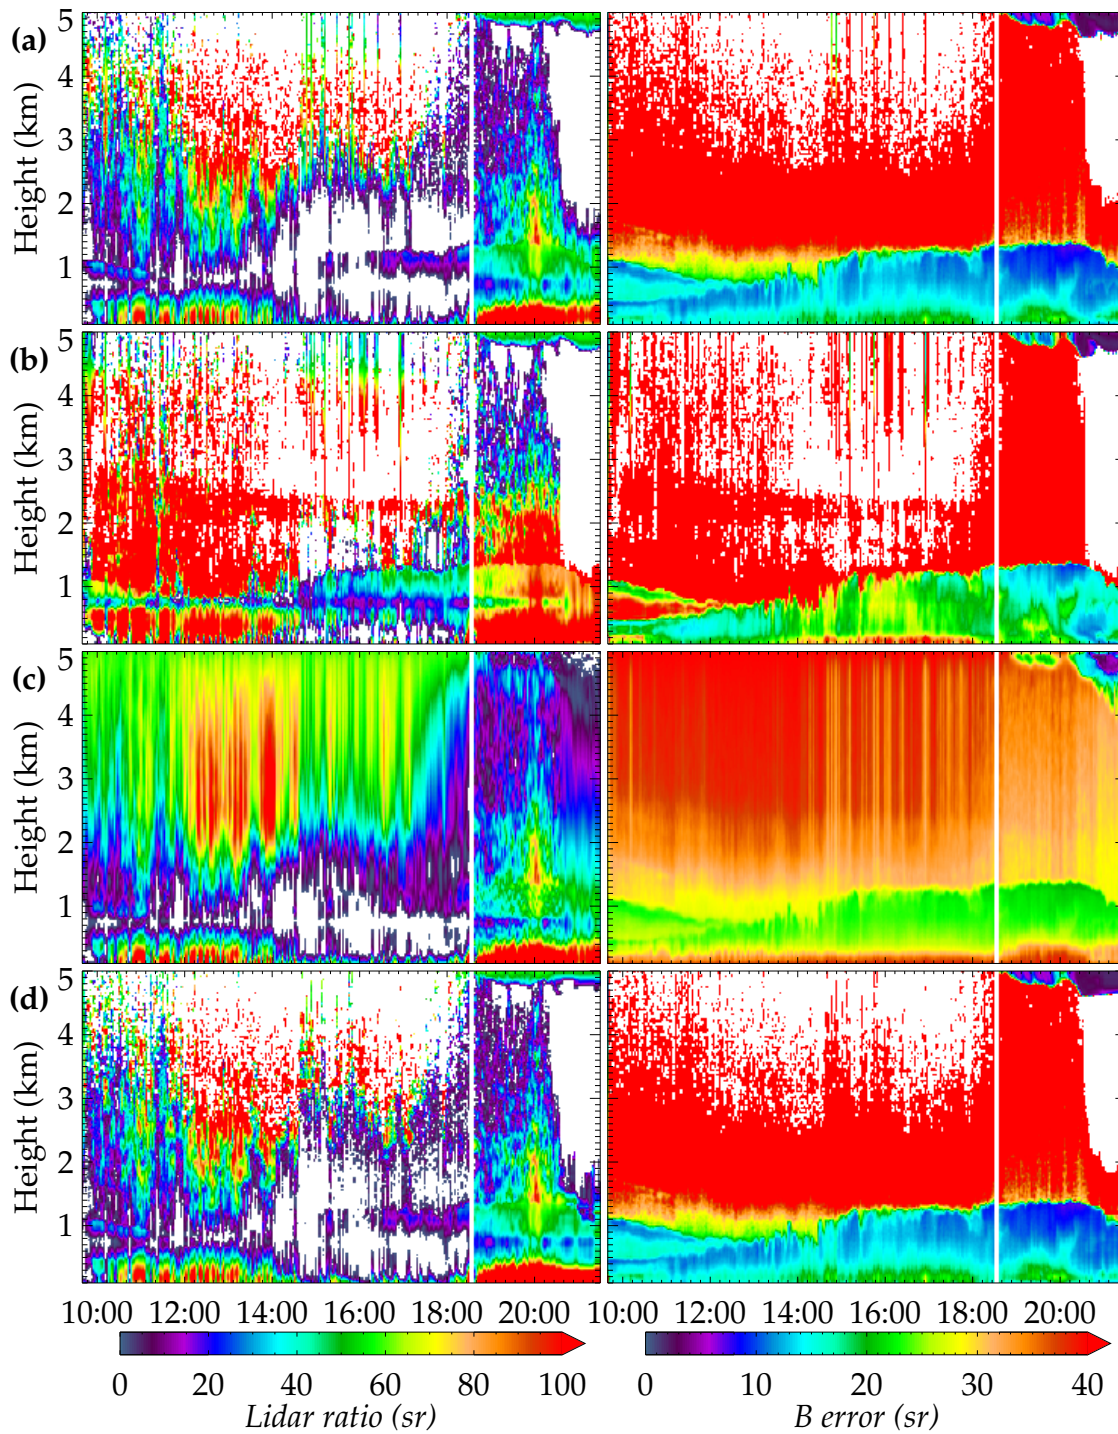


Figure 5.18: As fig. 5.17, but for the lidar ratio. Except in row (c), this is derived from α/β and their errors are propagated. White indicates missing data, $B = 1$ (the retrieval's lower bound), or $B = \infty$ (where $\beta = 0$).

However, the profile is not always bounded in the free troposphere, with β tending towards large values on a few occasions around 15:00. Also, the ingress of cloud at the top of the profile after 18:00 degrades the retrieval in the atmosphere beneath it. This was observed in Chapter 2 and remains troublesome but, as it does not affect values within the PBL, is most likely due to the *a priori* covariance. The bright scattering of a cloud is not included within the state space permitted to the retrieval by the *a priori* and so, though large values are selected, these are insufficient to fit the measurements. Vertical correlations allow the retrieval to compensate by minimising β beneath the cloud. The uncertainty shown here is misleading as, though the diagonal of S_x is small, the errors are greatly correlated. Further, the cost of these retrievals is enormous (order 10^3) and so the retrieval's ability to judge its own performance is severely impaired.

Row (b) of figure 5.17 shows the backscatter derived from analogue data. This field is qualitatively similar to the previous result, identifying similar features but roughly a factor of two smaller. It also exhibits greater contrast between high and low concentrations of aerosol than the PC data (especially around 10:00) and, as seen in fig. 5.16, indicates a cleaner troposphere than PC data¹⁵. The fractional error is similar to before, if slightly greater where errors in C and E_L dominate.

Row (c) returns to PC data but considers the log configuration of the retrieval. This is mostly identical to row (a), but the influence of β on the Raman channel has resulted in larger uncertainties throughout the PBL. The logarithmic retrieval has also selected greater scattering above 3 km and beneath cloud. The latter in particular is an undesirable response as it is less obvious that the retrieval has failed there.

Finally, row (d) uses each retrieval as the first guess for the next. This eliminates the unbounded solutions which were, in part, caused by the large first guess where there is no information available ($\beta = 10^{-5}$). The result is otherwise identical to row (a).

Figure 5.18 presents the lidar ratio fields for the same set of retrievals. These are more difficult to interpret as the lidar ratio is poorly constrained in the free troposphere, as indicated by the uncertainty reverting to its *a priori* value of 40 sr. The linear retrieval gives a speckled uncertainty field whilst the log configuration's is smooth because of the latter's greater vertical correlations and the former retrieving $\beta, \alpha = 0$. The results are distorted around 800 m by the incorrect N profile. This

¹⁵The information content is similar for the two measurements, so the exponential decay of the PC field is unlikely influenced by the *a priori*.

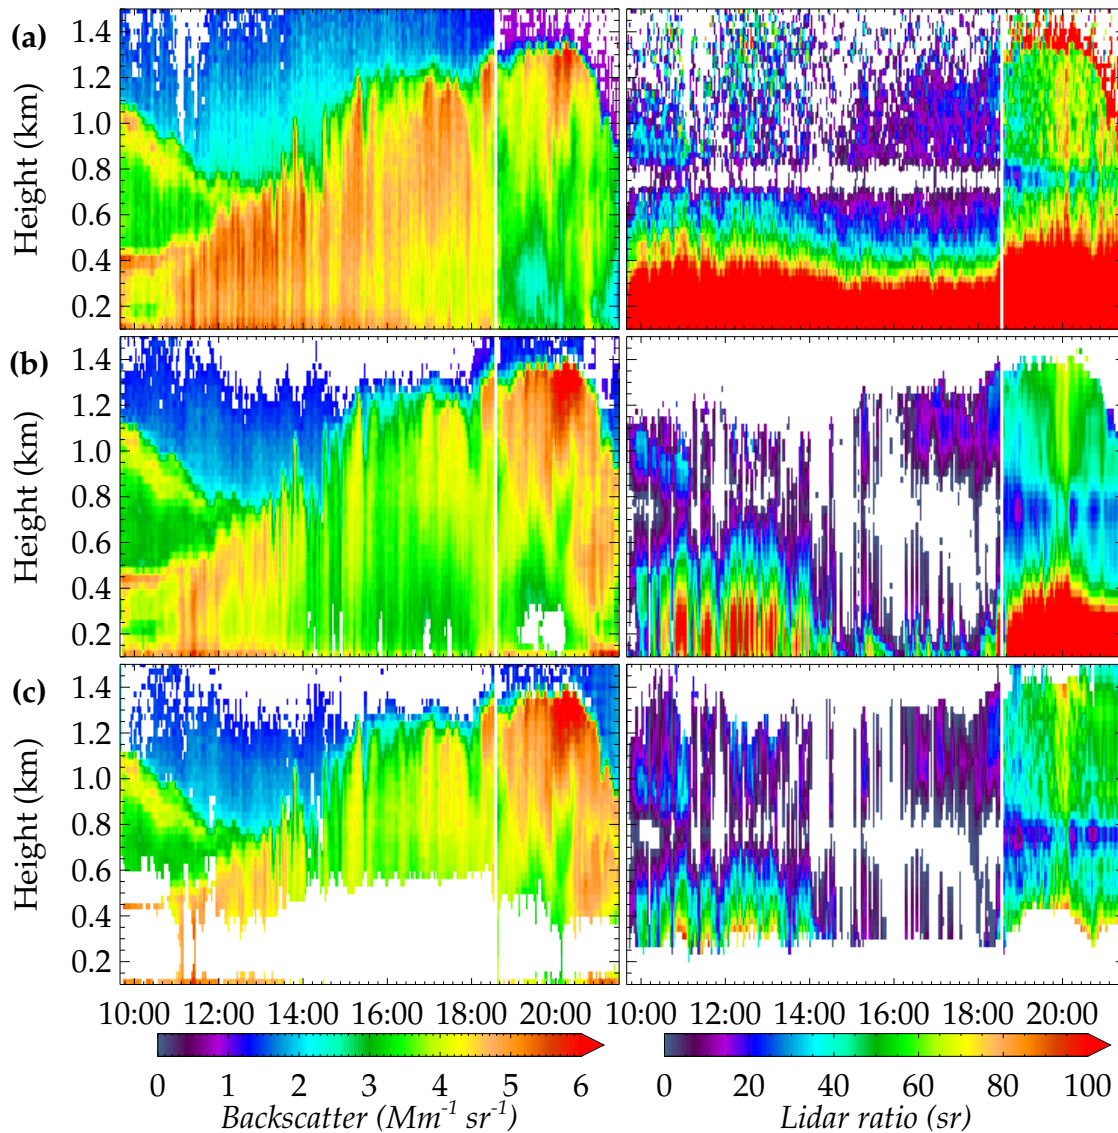


Figure 5.19: A closer examination of figures 5.17 and 5.18, filtering out data with an error greater than 20 % of β or 30 sr. **(a)** Results of the Ansmann method after averaging the data over 30 m to give a similar resolution to the retrieval. Only random errors are considered, estimated by Monte Carlo. **(b)** PC data, retrieving β and α . **(c)** PC data, retrieving $\ln \beta$ and B .

is less noticeable in the evening because the instrument was realigned¹⁶ at 18:30, increasing the magnitude of C . To compensate, B must increase to fit the brighter return.

If it is assumed that the alignment has drifted at 16:00, then the acceptable analogue results are entirely greater than 80 sr, which is excessively large. Therefore,

¹⁶The water vapour gain was adjusted for night-time observations and a slight change in alignment was noted in CFARR's log.




Figure 5.20: Additional observations from the BADC repository for 2 Mar 2010. **Left:** Depolarization ratio observed by the EZ lidar. **Right:** Vertical wind velocity observed by the Halo Doppler lidar.




Figure 5.21: Retrieved aerosol optical thickness at 355 nm (green: PC; black: analogue) compared to that observed by AERONET at level 2.0 (red) for 2 Mar 2010. The analogue Ansmann solution integrated between 0.25 and 2 km is shown in blue. **Inset:** Retrieved vs. AERONET χ_∞ , averaging all retrievals within 1 minute of an AERONET measurement, plotting the one-to-one line (dashed) and best fits (solid).

these observations will be neglected herein. To assess the PC results, figure 5.19 plots only data with an error less than 20 % or 30 sr (β and B respectively). Row (a) shows the application of the Ansmann method, where the data were averaged over 30 m to give a similar resolution to the retrieval. Rows (b) and (c) reproduce rows (a) and (c) of figures 5.17 and 5.18.

The three backscatter fields present very similar vertical profiles, but the retrievals and Ansmann method disagree by a factor of around 1.4 between 14:00 and 18:00 and 0.9 after then. This could be caused by an error in the magnitude of the calibration function or the retrieval of the laser energy, with the later considered more likely. The retrieval can compensate for such an error by scaling the backscatter profile, which would produce disagreements as observed. To resolve such a degeneracy in the forward model requires a more stringent *a priori* constraint than is currently provided. It is not clear what perturbed the system at 14:00, but the increase of the detector's gain at 18:00 would cause the observed decrease in backscatter.

The retrievals are consistent in their estimates of the lidar ratio and are no worse than the Ansmann method. The aerosol layer near 1 km at 10:00 gives a lidar ratio of around 30 sr. This is a residual layer lying above a developing mixed layer where lidar ratios are larger (around 50 sr). The low lidar ratio (and slightly lower depolarization ratio, see fig. 5.20) indicates large, mostly spherical particles, which is reasonable for an aged residual layer. The results are better in the evening, observing a peak B of 70 sr over a background of 50 sr, indicating the appearance of smaller particles. By this time, convective mixing has collapsed into a persistent updraft. The increase in depolarization ratio could indicate that newer, non-spherical particles are being lofted from the surface or are advected over the site. Advected aerosol is more likely considering the brevity of the peak.

Figure 5.21 compares χ_∞ deduced from either retrieval mode to AERONET measurements. PC observations give a slightly lower optical thickness (but within the predicted error) and neither appears to be well correlated with AERONET. However, the Ansmann algorithm gives equivalent values, indicating that the retrieval is giving the correct answer for the data it has been given - it is the magnitude of C that is wrong. There is also excessive variability over time, indicating the laser energy retrieval is wrong. This may result from the retrieval attempting to compensate for residual nonlinearity in the measurements with the laser energy.

It should be noted that all the retrievals have been presented without quality control. Figure 5.22 shows that, in the regions where the retrieval was performing




Figure 5.22: **Top:** Retrieval costs for 2 Mar 2010. **Bottom:** Laser energy as measured by a power meter compared to that retrieved, normalised by the value for text profile 6.

poorly, the cost of the retrieval is larger. The change in detector gain in the evening is evident from the large values of laser energy required to achieve a fit there.

5.1.5 Eyjafjallajökull ash

The eruptions of the Eyjafjallajökull volcano in southern Iceland during April and May of 2010 produced the single most significant volcanic ash event over northern Europe in the age of aviation [Schumann et al., 2011]. The closure of airspace cancelled around 100,000 flights [O'Dowd et al., 2012], inconveniencing millions of travellers across the globe and resulting in massive losses for airlines and related industries. Owing to the density of personnel and instrumentation within the reach of this plume, it has become one the most studied atmospheric events in history. The introduction of Johnson et al. [2012] provides a reasonable overview of the literature published to date and more will certainly be published over the years to come.

The CUV was operated, in addition to routine measurements, on April 19th to observe ash within the boundary layer. The optimal estimation retrieval was applied to these observations, assuming the calibration from test case 6; results are shown in figure 5.23. Plot (a) presents the depolarization ratio observed by the EZ

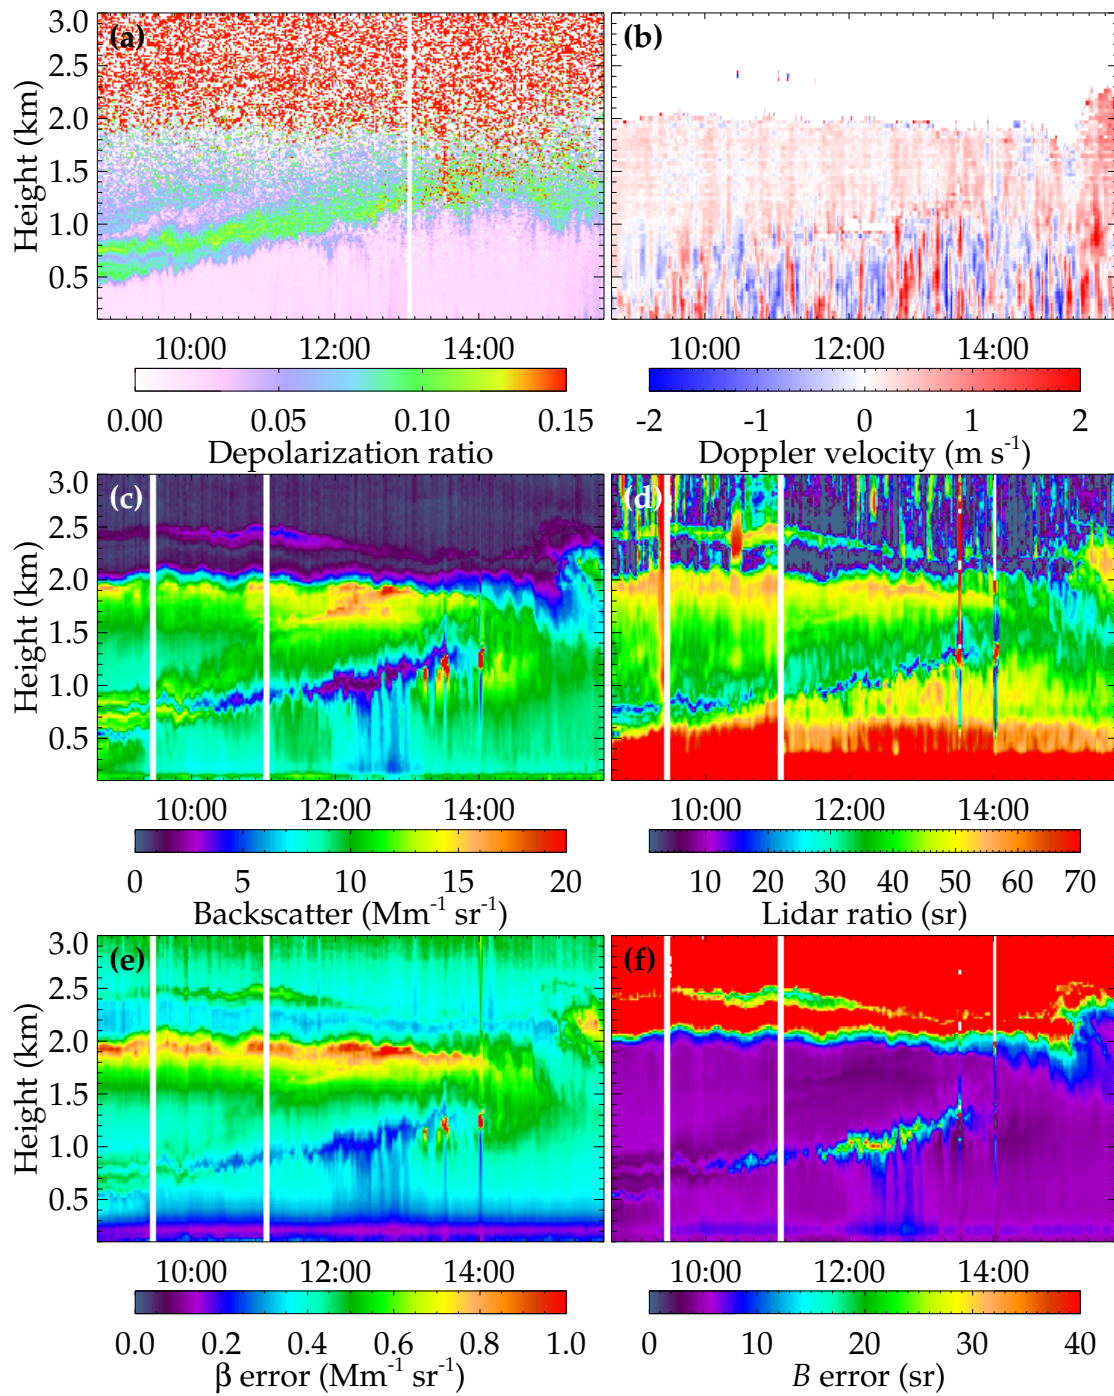


Figure 5.23: Observations of the Eyjafjallajökull ash plume at CFARR on 19 Apr 2010. **(a)** Depolarization ratio observed by the EZ lidar. Values above 1.9 km are dominated by noise. **(b)** Doppler (vertical) velocity observed by the Halo lidar. **(c)** Backscatter retrieved (in the linear mode) from CUV measurements. **(d)** Lidar ratio retrieved from same. **(e-f)** Retrieved errors on (c-d).




Figure 5.24: Distribution of extinction and backscatter for 19 Apr 2010. Lines delineate lidar ratios of 20, 40, 60, and 80 sr. **Left:** Points observed to have a depolarization ratio > 0.03 , which likely contain some quantity of ash. **Right:** The remaining points, corresponding to typical PBL aerosols. $\beta < 1 \text{ Mm}$ are observations in the free troposphere.

lidar at this time. Ash particles have a large depolarization due to their asphericity and the EZ measurements indicate the presence of a 400 m thick ash layer within the PBL, though the observed ratio of 0.1 is smaller than that expected for mineral dusts [0.35–0.37, Ansmann et al., 2011]. This aerosol remains distinct from the mixed layer, which grows over the morning (plot (b) highlights the increased turbulent mixing). Smaller depolarization observed above that layer may indicate that ash is mixed through the residual layer.

Plots (c) and (d) show the retrieved backscatter and lidar ratio. These indicate four separate aerosols:

- Within the mixed layer, $B = 40\text{--}60 \text{ sr}$ with minimal depolarization, broadly consistent with urban aerosols [Müller et al., 2007]. Backscatter is fairly homogeneous throughout this layer except during a period of updrafts around 13:00 when β decreases from 0.5 to $0.3 \text{ Mm}^{-1} \text{ sr}^{-1}$. The absence of a similar change elsewhere in the PBL gives some confidence that this is a real variation rather than a calibration artefact.
- A thin, highly depolarizing layer rests immediately above the mixed layer. It appears initially as bright, filamentary structures but becomes relatively dark

after 10:00. Raman water vapour observations indicate that the humidity here is half its value elsewhere in the PBL. The lidar ratio, though on the limit of resolution, is consistently lower than seen elsewhere in the PBL: 20–30 sr. This is well outside the ranges reported in literature for thin layers just above the PBL over southern England [82 sr, Marengo and Hogan, 2011]; in a thick, lofted layer over eastern Germany [50–65 sr, Ansmann et al., 2010]; or at the top of the PBL over France the previous day [52 sr, Hervo et al., 2012]. Combined with the smaller depolarization ratio, the retrieval indicates that the properties of the ash have changed significantly after 12–24 hours within the PBL¹⁷. The increase in B would imply a growth of the particles. Though the mixed layer is exhibiting significant variations, the ash remains distinct until the mixed layer reaches its full extent, indicating a capping temperature inversion or other feature isolates the air masses.

- A more weakly depolarizing aerosol resides in the residual layer. This persists until 14:00 when the mixed layer reaches its full extent and the three layers mix. This layer expresses much stronger backscatter and larger lidar ratios at its top, which then decrease with height. It is not entirely clear what mechanism could drive this. It may simply be stratification within a poorly mixed layer. Alternatively, smaller particles may have concentrated at the top of the layer whilst larger particles have begun to settle, though it does not seem that sufficient time has passed to produce so large a gradient.
- A thin layer of aerosol is present above the PBL. The EZ lidar did not resolve this, so no measure of the depolarization is available. Expressing lidar ratios of 40–60 sr with low β , the layer is consistent with aerosol typically observed at CFARR and there is no reason to suspect it is ash.

The variations are illustrated in figure 5.24, which presents the distribution of retrieved extinction and backscatter for all points with a depolarization ratio measured to better than 100 %. Lines of constant B are added for reference. Points likely to contain ash are highlighted by filtering for depolarizations greater than 0.03. In the left plot, the weakly depolarizing residual layer appears as a concentration of points around $\beta = 10 \text{ Mm}^{-1} \text{ sr}^{-1}$ and $B \simeq 40 \text{ sr}$. The mixed layer appears in the right plot as a more continuous distribution between $B = 40$ and 60 sr. The failure of the retrieval near the surface is evident in the vertical spread of points at

¹⁷And this occurs without the influence of water.

$\beta = 8 \text{ Mm}^{-1} \text{ sr}^{-1}$. In the free troposphere, $\beta < 1 \text{ Mm}^{-1} \text{ sr}^{-1}$, where poorly constrained retrievals produce a broad distribution in both plots. There are few observations of the thin ash layer, but their presence is evident from the presence of observations near $B \simeq 20 \text{ sr}$ in the left plot not expressed on the right.


5.2 RACHEL

RACHEL was stationed at CFARR from January to May of 2010 during which it was operated for a total of 471 hours over 33 days, including continuous observations covering 12–19 March and 15–25 April. As the air conditioner provided insufficient cooling, it was necessary to ventilate the instrument every 30 minutes on April afternoons. The inability to properly cool the laser had an increasingly detrimental effect on its stability, with spontaneous and frequent changes in energy of up to a factor of two. Also, a faulty optical mounting required daily realignment, such that the continuous presence of an operator was necessary. It became evident that the calibration function varied quite substantially between these realignments, which inspired the work of Chapter 4.

After this campaign, various repairs were made by Hovemere to improve the air conditioner and scanning system (which had not been used at CFARR due to issues with the motors). An accident required the replacement of the scanning mirrors and the laser had to be returned to its manufacturer for several months, after which its stability vastly improved. The system was then installed (by crane) on the roof of the Atmospheric Physics Building in Oxford (51.7596° , -1.2557° , 14 amsl) from June 2011, including the addition of a new laser interlock system.

Operations were problematic from the beginning. The waterproofing of the instrument was found to be insufficient for an exposed position, with leaks destroying two computers. Several months were required to source suitable replacements, though it was never possible to interface the scanning system with the new hardware, such that scanning operation was abandoned. Further, conflicts between the hardware and software caused unpredictable computer crashes. Several weeks were spent writing new control scripts in LabView to eliminate this issue. Despite extensive effort, the instrument could never be fully sealed such that the main window would become covered in mist on chill mornings, completely obscuring any signal and hindering observations during winter. The alterations to the air conditioner were insufficient as temperature control remained a problem during the

Figure 5.25: Retrieved overlap functions for RACHEL from 15 Apr 2010. The purple-blue line is retrieved from an hour of observation from 17:00, compared to the Wandinger solution in red. The orange-green line derives from two hours after 18:00, compared to the black line.



summer, making operation impossible in direct sunlight. Finally, most components of the laser failed at one time or another, most likely due to the unfavorable conditions. It could not be removed from the roof due to its size and weight, such that even minor repairs required several additional hours. By the end of 2012, work on RACHEL was ceased as beyond reasonable efforts to repair in such conditions.

Overall, 710 hours of data were collected from 80 days over the two years. Much of this is meaningless due to failures of the acquisition software, a loss of laser energy, drifting alignment, or misting of the main window. The remainder is difficult to analyze due to a lack of suitable measurements to determine $C(R)$ or τ_d . It was decided not to include data from the Oxford campaign in this thesis as these terms would introduce excessive uncertainty to any retrievals made. Though the CFARR data are similarly afflicted, there is auxiliary information available to evaluate the calibration and retrievals. Further, the observations of Eyjafjallajökull ash between 15–25 April were of scientific interest. A selection of those results are presented in this section.

The vertical offset was found using the technique of Section 5.1.1 to be $-27(8)$ m. The dead time corrections were estimated to be $36(4)$ and $44(4)$ ns in Chapter 3. Figure 5.25 presents the retrieval of C from profiles averaged over one and two hours on 15 Apr 2010 (starting at 17:00 and 18:00, respectively). These test the lim-




Figure 5.26: As fig. 5.23, but as observed by RACHEL over Apr 15-16 2010.




Figure 5.27: As fig. 5.26, but over Apr 17-19 2010.




Figure 5.28: As fig. 5.24, but for RACHEL observations on 16 Apr 2010.

its of the Wandering algorithm, which is unable to converge. The solution for the two-hour profile diverges slowly, with the profile shown (black) being the result after 15 iterations. If evaluation were continued indefinitely, it would resemble the shorter observation (red). This is likely the source of the discrepancies between that scheme and the retrieval. Otherwise, the retrieval appears to be performing excellently. It is hoped that the calibration does not drift as significantly as these two profiles imply and that their differences are due to insufficient information in the shorter profile.

RACHEL was operational from 14:30 on April 15th till 12:00 on April 19th when electrical work required a shut down. The entire period was processed at a resolution of 33 m and 10 minutes, as shown in figures 5.26 and 5.27. These are impressive as, despite overwhelming difficulties in the calibration, the retrieval has converged to profiles within the realm of reality, which is more than other methods of analysis could achieve. For the first two days (fig 5.26):

- The retrieval still performs poorly in the presence of cloud, with very low β beneath clouds. It is not affected in profiles between clouds.
- The laser energy is not well retrieved, producing sharp changes in β from 2 to 5 $\text{Mm}^{-1} \text{sr}$.
- The retrieval is not always properly converged, as demonstrated by the step decrease in scattering within a cloud at 7:00 on the 16th.

- The alignment is, as noted at the time, drifting. The elastic channel appears to be affected in late-afternoon (see 18:00 on the 16th) and the Raman channel overnight. The lidar ratio is, for that reason, mostly meaningless.

The appearance of ash over CFARR after 12:00 on the 16th is retrieved despite the presence of clouds beneath. Within the plume, $B = 50\text{--}60\text{ sr}$ is observed, in much closer agreement with independent observations. The quoted uncertainty of 20 sr is large, and is certainly an underestimate, but impressive considering the quality of the data. Filtering retrievals by the observed depolarization, figure 5.28 highlights a concentration of measurements of $50 < B < 60\text{ sr}$ for depolarizing aerosol whilst there is no consistent lidar ratio for spherical aerosol¹⁸. This implies that the ash, though mixed into the PBL, remains distinct over a period of several hours and apparently requires several days to evolve towards the smaller B observed by the CUV.

Unfortunately, the alignment drifts substantially on the night of the 16th such that the remaining retrievals (fig 5.27) are useless for comparison to the CUV observations. Calibration could be attempted at night, when the aerosol appears to be stably stratified, but χ_∞ will be large and uncertain. In the morning, when AERONET measurements are available, the aerosol distribution is too inhomogeneous to trust any calibration that is returned.

¹⁸There is a slightly greater concentration in the range $60\text{--}80\text{ sr}$.

Chapter 6

Conclusions

When the mechanisms behind an observation are sufficiently well understood that they can be described mathematically, it is possible to use optimal estimation retrieval to determine the conditions statistically most-likely to produce that observation [Rodgers, 2000]. An example relevant to the atmospheric sciences would be the equations of radiative transfer relating the number and size of droplets in a cloud to the radiance from them observed at several wavelengths by a satellite, given some pressure and temperature information.

Rather than apply artificial constraints such as smoothness to better constrain the problem, optimal estimation weights its solution against any *a priori* knowledge that may be available. Physical processes, such as wind-driven mixing or synoptic-scale conditions, will introduce correlations between observations over some region of space. Another example could be an aerosol plume that is known from back-trajectory analysis to be comprised of desert dust, for which the size distribution and refractive index of the particles can be constrained. Such information allows the retrieval to make the most efficient use of the data it is given, using physically-consistent relationships to estimate the desired quantities without the need to impose *ad hoc* corrections or transformations to the measurements.

The result of such retrievals will be a combination of the information provided by the measurements and the *a priori* assumptions such that, relative to more traditional techniques, there will be a distortion towards solutions consistent with the *a priori*. For some researchers, this is a disadvantage of the technique — introducing ‘biases’ to the products. Such criticism misses several important points. Firstly, the retrieval provides several metrics (the covariance matrix and averaging kernel) to assess the impact of the *a priori*. Results dominated by these assumptions will express greater uncertainty and will have minimal information content. Equivalent results produced by traditional techniques will exhibit a greater uncertainty

but provide no indication of the information content of that result. That is left to quality control analyses, which often include some manner of judgement against expected results. Clearly traditional analyses are not hindered by this limitation, but (from a statistical standpoint) that judgement is important to the result. Optimal estimation provides a more consistent means to introduce such judgements (ubiquitous throughout science) into data and error analysis.

Secondly, it does not seem erroneous for a retrieval, when there is incomplete information available, to make use of prior knowledge. In the modern day, no observation exists in a vacuum. Assessing the accuracy and relevance of the plethora of existing information, in addition to properly introducing it into the analysis, is a significant challenge but is absolutely necessary to apply, understand, and validate the new observations. In the work presented, this is probably best illustrated in figure 4.24, comparing the performance of an optimal estimation retrieval of the calibration function to the technique of Wandinger and Ansmann [2002]. For these data, the information content decreases with height. The iterative solution responds to this by producing increasingly disperse values above 2 km, which are consistent with the data but not actually useful for estimating C . With optimal estimation, the forward model forces the result to vary smoothly, consistent with geometric optics, and uses information lower in the profile to deduce answers that are (in most cases) more accurate than existing techniques.

Optimal estimation is not an ideal technique; it can only find solutions within the expected range. In some circumstances this is not a great limitation, such as case 13 of figure 4.24, where the retrieval could not properly fit features in C caused by components beyond its forward model but still obtained a solution of comparable accuracy to existing techniques. In other cases, this limitation completely eliminated the retrieval's utility, such as the inability of the scheme outlined in Chapter 2 to process data containing cloud. Its *a priori* does not permit values consistent with cloud particles and so the retrieval distorts the remainder of the profile in an attempt to fit the data using permitted variations. In such cases, the cost of the retrieval will be large and reveal the mismatch between reality and the assumed model. If information were required there, either a new, less-stringent model or a different analysis technique entirely could be applied.

Such limitations are the most obvious criticism of the retrieval scheme developed in Chapter 2. There, the widely-known lidar equations (2.2.10)–(2.2.12) were used as a forward model to retrieve aerosol backscatter and either extinction or the lidar ratio from two-channel Raman lidar data. Though it is possible within such a

retrieval to estimate microphysical properties of the aerosol directly, the assumptions necessary to do so are poorly constrained. It was deemed more useful to concentrate on the backscatter, extinction, and lidar ratio (as widely derived in the lidar community). Extinction and backscatter are roughly proportional to aerosol number density whilst their ratio is inversely proportional to particle size. Such observations provide a qualitative understanding of aerosol abundance and its evolution with time and have been successfully used to evaluate PBL height for numerical weather prediction, as a validation of satellite cloud top height retrievals, and in the monitoring of urban pollution. However, lidar has yet to make as significant a contribution in the study of aerosol dynamics and its impacts on climate as it should have considering its superior temporal and vertical resolution.

The proposed retrieval scheme may be a first step towards a more widespread understanding and use of lidar data. In the analysis of simulated and real data, the proposed retrieval is consistent with existing analyses. In addition, it returns a more rigorous estimate of the uncertainty of its profiles and selects the resolution most appropriate to the quality of data. Backscatter was always retrieved at the finest resolution allowed (mostly 33 m, but this remains true at the instrumental limit < 10 m) and with an uncertainty between 2 % in ideal circumstances and 20 % in the least. Extinction and the lidar ratio are less well constrained, expressing resolutions of 300–500 m in simulations and 0.1–1 km with real data. Importantly, these are different from the scale of vertical correlations assumed *a priori* and increase with height. The retrieval is selecting the most suitable resolution independently, unlike the smoothing filters used in most studies [e.g. Turner et al., 2002]. The uncertainty in extinction retrieved from real data is fairly large (> 15 %), but that is in part due to the short time scales evaluated (one minute in figs. 5.18 and 5.23). The integration time can be increased to reduce these errors to any desired level (at least until atmospheric variability begins to dominate) and they are comparable to, if not smaller than, fair estimates of the error resulting from the standard Raman lidar technique of Ansmann et al. [1992] applied to the same data, as shown in figure 5.12.

The magnitude of the uncertainty was shown to be dominated by the calibration of the instrument — primarily the offset of its vertical axis, the nonlinear response of its detectors, and the calibration function. Ideally, all of these would be determined with dedicated laboratory measurements, especially the offset and nonlinearity which should change very little over time. However, the data available were either collected by an independent group with little need to consider

such corrections or at a very early stage in a student's studies, when the importance of thorough calibration had not yet become evident. Though the data could simply have been disregarded, the aim was to extract as much information as possible and test the limits of lidar data analysis.

Little success was made in the estimation of the nonlinear response of a PMT. It can be measured by collecting two profiles in succession — one attenuated by a neutral density filter and another at full strength. Most commonly, the detector is assumed to be 'paralysed' for a fixed time after any detection, giving a simple analytic correction (1.3.23). This model was found to be useful for a variety of more realistic responses but over a limited dynamic range. As such, regression techniques failed as they could not judge which data were well-modelled and which were not in the fitting. A technique proposed by Whiteman et al. [1992] was found to be the most suitable means of fitting such a correction to data as a scientist could more easily make this distinction. Attempts to fit more complex nonlinear corrections to data failed as lidar profiles do not provide sufficient information to distinguish between the different models.

Fortunately, that was not the most significant source of uncertainty. A myriad of techniques have been proposed to determine the overlap function and calibration constant of a lidar from a variety of dedicated measurements. The utility of such schemes is not in doubt, but most make use of a unique instrumental design or specific meteorological conditions to better constrain the problem and so cannot be applied to data *a posteriori* nor transferred to a different instrument. As the overlap function will drift over time due to the fundamental limitations of a laser and other optical components, it is preferable to determine the overlap function frequently and from routine measurements.

Chapter 4 proposed use of the formulation of Halldórsson and Langerholc [1978] in conjunction with an idealised aerosol profile (an error function, as posited by Steyn et al. [1999]) as a forward model to estimate the calibration function (the product of the overlap function with a calibration constant) from observations of a reasonably well-mixed PBL by a Raman channel and a measure of the aerosol optical thickness. An appropriate configuration for the model was selected by considering the retrieval's ability to reproduce simulations while demonstrating minimal sensitivity to its parameters, such as the height range considered. The *a priori* was based on the OPAC model and measurements of various components.

This retrieval successfully estimated idealised calibrations to within a few percent, even in the presence of significant perturbations to the simplified aerosol dis-

tribution. It had some difficulty reproducing C for instruments that did not conform to the assumed optical system, such a laser with a Gaussian beam shape. This could be avoided by annealing the retrieval, whereby the retrieval is repeated a few dozen times with randomised first guesses. The solution among those with the lowest cost was found to be in very good agreement with that simulated provided the retrieval selected a realistic aerosol profile. When applied to real observations, results were found to be consistent with those of the algorithm proposed by Wandinger and Ansmann [2002]. Total uncertainties were 5–15 %. Discrepancies could be attributed to differences between the overlap functions of the elastic and Raman channels, as expected in any lidar system due to inhomogeneities in the optical components.

Hypothetically, the retrieval scheme proposed could also be applied to the elastic channel to determine the overlaps separately (rather than assuming they are equal, as done in Chapter 5). This would demonstrate significantly greater uncertainties as the backscatter has an identical impact on the elastic profile as the calibration, such that its variations will be more important than in the scheme outlined (which is only sensitive to the optical thickness, an integrated quantity). It is, though, an obvious next step in the examination of lidar calibration.

The various retrievals were then applied to a selection of observations by the CUV and RACHEL at CFARR during 2010. These demonstrated varying degrees of success, but most failures could be attributed to limitations of the external information. In figure 5.17, a discontinuity in N imprecisely captured by radiosonde distorts the retrieval. Observations of the PBL by RACHEL are contorted by rapid changes in its alignment. In each case, other techniques suffer identical failures. Remembering these difficulties (but not thoroughly accounting for them as that eliminates any information content in the observations), the retrieval was applied to several days of observation of ash from the Eyjafjallajökull eruption of April 2010. Lidar ratios of 50–60 sr were observed when it first appeared over CFARR, which persisted as it was mixed into the PBL. Several days later, this was observed to have decreased to 20–30 sr within a distinct layer of highly depolarizing particles and to around 40 sr for less depolarizing particles mixed into a residual layer. If the retrieval is to be believed, these indicate that the ash particles grew over a period of several days (or at least were altered in some significant manner).

It has been shown that the application of optimal estimation produces estimates of the aerosol scattering no worse those currently produced with Raman lidar, but with additional information that can be of great use to those unfamiliar with the

vagaries of lidar operation. It also provides an automated means to adjust the resolution of extinction products as the SNR of the data decreases. The next stage in the development of the extinction/backscatter retrieval is to improve its performance in the presence of cloud. This requires an alteration of the *a priori* assumptions to accommodate the divergence between the aerosol and cloud layers, which may be possible through a binary flag (whereby once a bin exceeds some threshold of β , it is allotted different properties). A consideration of multiple scattering may also be necessary, such as that outlined in Eloranta [1998]. For the instruments considered, this should be a small correction but will be significant within clouds or for space-based systems. As the lidar ratio configuration of the retrieval was found to be inferior in theory and in practice, it would also be advisable to further research the *a priori* correlation of α and β . The current value appears to be important only in its presence rather than magnitude, but there must be some impact on the retrieval. Once more data have been processed with this scheme, it should be possible to consider the correlations present and produce a new covariance matrix.

The greatest source of uncertainty in lidar analysis has been shown to be the calibration of the instrument and there remains significant work to be done in reducing this. The calibration constant is poorly estimated by all methods attempted. Future retrievals may require a more stringent constraint than aerosol optical thickness. Also, the current forward model tends to retrieve decreasing C as an result of fitting $\phi_{\parallel,\perp}$ to features near the surface. As the most prevalent *a priori* assumption in any estimation of the overlap function is that it is constant with height above some point, this demands further attention.

The proposed retrieval indicates that there is no issue with attempting to retrieve aerosol in the region where the overlap function is rapidly varying, provided there is an reasonable assessment of the uncertainty in its estimation. However, it is unlikely that such an assessment could be made. Estimating the calibration function requires a knowledge of the aerosol present and this will be most uncertain and variable in the surface layer immediately above the instrument. The use of an optical system to compress the physical distances involved can eliminate the impact of aerosol but will distort the region of most rapid change. As such, the uncertainties in the first few hundred metres of a lidar profile will be excessive. Though retrieval is possible, it will provide minimal useful information.

Otherwise, the most obvious direction for future work is to simply apply the retrieval schemes to more data, especially instruments that have been thoroughly calibrated in a laboratory setting. Minimisation of the extraneous uncertainties

will provide the most convincing demonstration of the application of optimal estimation retrieval in the analysis of lidar observations.

Appendix A

Additional derivations

A.1 Lidar equation Jacobean matrix

The Jacobean, $\mathbf{K} = \partial \mathbf{F} / \partial \mathbf{x}$, must be determined with each evaluation of the forward model as it is a function of \mathbf{x} . It can be determined numerically, but analytic expressions are preferable as they should be less computationally expensive and minimise errors. Though not trivial, the forward model of Section 2.2 does have analytic derivatives.

In the equations that follow, if the el/ra superscript is neglected, the expression holds for either. Considering the dead time correction (2.2.10),

$$\frac{\partial \varphi_i}{\partial x} = \frac{\partial \varphi_i}{\partial E_i} \frac{\partial E_i}{\partial x} = \frac{M E_i}{(1 + \frac{\tau_d}{\tau_b} E_i)^2} \frac{\partial E_i}{\partial x}. \quad (\text{A.1.1})$$

A.1.1 Extinction configuration

Considering equations (2.2.11) and (2.2.12) where β and α are independent,

$$\frac{\partial E_i^{\text{el}}}{\partial \beta_k^r} = E_L \frac{C_i^{\text{el}}}{R_i^2} \frac{\partial \beta_i}{\partial \beta_k^r} \exp \left[-2(\sigma_R^{\text{el}} \mathcal{N}_i + \chi_i) \right] \quad (\text{A.1.2})$$

$$\frac{\partial E_i^{\text{el}}}{\partial \alpha_k^r} = -2(E_i^{\text{el}} - E_B^{\text{el}}) \frac{\partial \chi_i}{\partial \alpha_k^r} \quad (\text{A.1.3})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial \beta_k^r} = 0 \quad (\text{A.1.4})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial \alpha_k^r} = -(E_i^{\text{ra}} - E_B^{\text{ra}}) \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \frac{\partial \chi_i}{\partial \alpha_k^r}. \quad (\text{A.1.5})$$

To derive $\frac{\partial \beta_i}{\partial \beta_k^r}$, assume \mathbf{E} is invertible such that (2.2.8) gives,

$$(\beta_k^r)'' = \sum_{l=1}^{n-2} (\mathbf{E}^{-1} \mathbf{F})_{kl} \beta_l^r \quad (\text{A.1.6})$$

$$\therefore \frac{\partial(\beta_k^r)''}{\partial\beta_l^r} = \begin{cases} (\mathbf{E}^{-1}\mathbf{F})_{kl} & \text{if } 1 \leq k \leq n-2, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.1.7})$$

Thus, differentiating (2.2.7),

$$\frac{\partial\beta_i}{\partial\beta_k^r} = A_{ij}\delta_{jk} + B_{ij}\delta_{(j+1)k} + C_{ij}(\mathbf{E}^{-1}\mathbf{F})_{jk}(1 - \delta_{k0}) + D_{ij}(\mathbf{E}^{-1}\mathbf{F})_{(j+1)k}[1 - \delta_{k(n-1)}], \quad (\text{A.1.8})$$

where δ_{jk} is the Kronecker delta.

It is then known from the total derivative that,

$$\frac{\partial\chi_i}{\partial\alpha_k^r} = \sum_{l=0}^{n-1} \frac{\partial\chi_i}{\partial\chi_l^r} \frac{d\chi_l^r}{d\alpha_k^r}. \quad (\text{A.1.9})$$

Considering (2.2.15),

$$\frac{d\chi_l^r}{d\alpha_k^r} = \begin{cases} 0 & l < k, \\ \frac{1}{2}(r_k - r_{k-1}) & l = (k \neq 0), \\ \frac{1}{2}(r_{k+1} - r_{k-1}) & l > (k \neq 0), \\ \frac{1}{2}(r_1 + r_0) & l > (k = 0), \\ r_0 & l = k = 0. \end{cases} \quad (\text{A.1.10})$$

As identical vertical grids are used,

$$\frac{\partial\chi_i}{\partial\chi_l^r} = \frac{\partial\beta_i}{\partial\beta_l^r}. \quad (\text{A.1.11})$$

Hence,

$$\frac{\partial\chi_i}{\partial\alpha_k^r} = \begin{cases} r_0 \frac{\partial\beta_i}{\partial\beta_0^r} + \frac{r_1 + r_0}{2} \sum_{l=1}^{n-1} \frac{\partial\beta_i}{\partial\beta_l^r} & k = 0, \\ \frac{r_k - r_{k-1}}{2} \frac{\partial\beta_i}{\partial\beta_k^r} + \frac{r_{k+1} - r_{k-1}}{2} \sum_{l=k+1}^{n-1} \frac{\partial\beta_i}{\partial\beta_l^r} & \text{otherwise.} \end{cases} \quad (\text{A.1.12})$$

A.1.2 Lidar ratio configuration

When β and α are dependent,

$$\frac{\partial E_i^{\text{el}}}{\partial\beta_k^r} = E_L \frac{C_i^{\text{el}}}{R_i^2} \frac{\partial\beta_i}{\partial\beta_k^r} \exp[-2(\sigma_R^{\text{el}}\mathcal{N}_i + \chi_i)] - 2(E_i^{\text{el}} - E_B^{\text{el}}) \frac{\partial\chi_i}{\partial\beta_k^r} \quad (\text{A.1.13})$$

$$\frac{\partial E_i^{\text{el}}}{\partial B_k^r} = -2(E_i^{\text{el}} - E_B^{\text{el}}) \frac{\partial\chi_i}{\partial B_k^r} \quad (\text{A.1.14})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial \beta_k^r} = -(E_i^{\text{ra}} - E_B^{\text{ra}}) \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \frac{\partial \chi_i}{\partial \beta_k^r} \quad (\text{A.1.15})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial B_k^r} = -(E_i^{\text{ra}} - E_B^{\text{ra}}) \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \frac{\partial \chi_i}{\partial B_k^r}, \quad (\text{A.1.16})$$

where,

$$\frac{\partial \chi_i}{\partial \beta_k^r} = B_k^r \frac{\partial \chi_i}{\partial \alpha_k^r}, \quad \frac{\partial \chi_i}{\partial B_k^r} = \beta_k^r \frac{\partial \chi_i}{\partial \alpha_k^r}.$$

A.2 Lidar equation parameter errors

A.2.1 Dead time, τ_d

Without comment on forward model errors due to the representation of the dead time correction,

$$\frac{\partial \varphi_i}{\partial \tau_d} = -\frac{\varphi_i^2}{M\tau_b}. \quad (\text{A.2.1})$$

For the remaining terms,

$$\frac{\partial \varphi_i}{\partial b} = \frac{\varphi_i}{1 + \frac{\tau_d}{\tau_b} E_i} \frac{\partial E_i}{\partial b}. \quad (\text{A.2.2})$$

A.2.2 Background counts, E_B

$$\frac{\partial E_i}{\partial E_B} = 1. \quad (\text{A.2.3})$$

A.2.3 Calibration function, C

$$\frac{\partial E_i}{\partial C_k} = \delta_{ik} \frac{E_i - E_B}{C_i}. \quad (\text{A.2.4})$$

A.2.4 Vertical offset, R_0

$$\begin{aligned} \frac{\partial \text{spline}_{r \rightarrow R_i} x^r}{\partial R_0} &= \frac{x_{j+1}^r - x_j^r}{r_{j+1} - r_j} \\ &+ \frac{[(r_{j+1}r_j)^2 - 3(r_{j+1} - R_i)^2](x_j^r)'' + [3(R_i - r_j)^2 - (r_{j+1} - r_j)^2](x_{j+1}^r)''}{6(r_{j+1} - r_j)}. \end{aligned} \quad (\text{A.2.5})$$

Hence,

$$\frac{\partial E_i^{\text{el}}}{\partial R_0} = -2E_i^{\text{el}} \left(\frac{1}{R} + \frac{\partial \chi_i}{\partial R_0} \right) + E_L \frac{C_i^{\text{el}}}{R_i^2} \exp[-2(\sigma_R^{\text{el}} \mathcal{N} + \chi_i)] \frac{\partial \beta_i}{\partial R_0} \quad (\text{A.2.6})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial R_0} = -E_i^{\text{el}} \left[\frac{2}{R} + \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \frac{\partial \chi_i}{\partial R_0} \right], \quad (\text{A.2.7})$$

which neglects any dependence of N or χ^{m} on R_0 due to interpolation.

A.2.5 Ångstrom coefficient, κ

The Ångstrom coefficient is assumed equal to unity (2.1.14). This is unlikely for any particular aerosol, though the impact of this assumption was shown to be negligible in figure 2.19.

$$\frac{\partial E_i^{\text{el}}}{\partial \kappa} = 0 \quad \frac{\partial E_i^{\text{ra}}}{\partial \kappa} = -\chi_i (E_i^{\text{ra}} - E_B^{\text{ra}}) \left(\frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right)^{\kappa} \log \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}}. \quad (\text{A.2.8})$$

A.2.6 Number density, N

The error on N should include interpolation errors for radiosonde data or natural variations about the standard atmosphere in addition to any measurement error.

$$\frac{\partial E_i^{\text{el}}}{\partial N_k} = \delta_{ik} E_L \frac{C_i^{\text{el}}}{R_i^2} \frac{\sigma_R^{\text{el}}}{\mathcal{B}} \exp [-2(\sigma_R^{\text{el}} \mathcal{N}_i + \chi_i)] \quad (\text{A.2.9})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial N_k} = -\delta_{ik} \frac{E_i^{\text{ra}}}{N_i}. \quad (\text{A.2.10})$$

A.2.7 Molecular optical thickness, χ^{m}

Though, strictly, this is dependent on N , it is interpolated independently or calculated assuming a hydrostatically-balanced atmosphere. This relation will be ignored both here and in \mathbf{S}_b .

$$\frac{\partial E_i^{\text{el}}}{\partial \chi_k^{\text{m}}} = -2 \delta_{ik} \sigma_R^{\text{el}} (E_i^{\text{el}} - E_B^{\text{el}}) \quad (\text{A.2.11})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial \chi_k^{\text{m}}} = -\delta_{ik} (\sigma_R^{\text{el}} + \sigma_R^{\text{ra}}) (E_i^{\text{ra}} - E_B^{\text{ra}}). \quad (\text{A.2.12})$$

A.3 Calibration function Jacobean matrix

Due to the use of Rhomberg integration in the calibration forward model, it is both impractical to determine the Jacobean numerically and inaccurate to use analytic expressions due to errors in the integration. The case structure of S_H ensures that estimation of \mathbf{K} by finite differences is further complicated. As such, particular care must be taken to derive accurate, practical expressions.

As for the previous forward model,

$$\frac{\partial \varphi_i^{\text{ra}}}{\partial x_k} = \frac{\varphi_i^{\text{ra}}}{1 + \frac{\tau_d}{\tau_b} E_i^{\text{ra}}} \frac{\partial E_i^{\text{ra}}}{\partial x_k}. \quad (\text{A.3.1})$$

The first ten elements of the state vector only affect the effective area such that,

$$\frac{\partial E_i^{\text{ra}}}{\partial x} = \frac{E_i^{\text{ra}} - E_B^{\text{ra}}}{A(R_i, \epsilon)} \frac{\partial A}{\partial x}. \quad (\text{A.3.2})$$

A.3.1 Axis separation, d

Three parameters only enter A through d such that,

$$\frac{\partial A}{\partial \delta} = \frac{\delta + \phi_{\parallel} R}{d} \frac{\partial A}{\partial d} \quad \frac{\partial A}{\partial \phi_{\parallel}} = R \frac{\delta + \phi_{\parallel} R}{d} \frac{\partial A}{\partial d} \quad \frac{\partial A}{\partial \phi_{\perp}} = \frac{\phi_{\perp} R^2}{d} \frac{\partial A}{\partial d}, \quad (\text{A.3.3})$$

where,

$$\frac{\partial A}{\partial d} = \left(\frac{\gamma}{\nu w} \right)^2 \left[\left| \frac{\partial S_H}{\partial d} \right|_{\rho_T} - \left| \frac{\partial S_H}{\partial d} \right|_{\rho_o} \right]. \quad (\text{A.3.4})$$

The derivative is clearly zero for $|w - d| \geq a + \rho$ or $w + d \leq |a - \rho|$. Then considering (A.7.5),

$$\frac{\partial S_H}{\partial d} = \begin{cases} 0 & |w - d| \geq a + \rho, \\ & \text{or } w + d \leq |a - \rho|, \\ - \int_{\max[|w-d|, |a-\rho|]}^{\min[w+d, a+\rho]} \frac{\sqrt{\Upsilon(\mu, w, d) \Upsilon(a, \rho, \mu)}}{\pi \mu d} d\mu & \text{otherwise.} \end{cases} \quad (\text{A.3.5})$$

Though it is also possible to derive a similar expression from (A.7.18), that has a pole at the limits of the integral which complicates numerical evaluation.

A.3.2 Beam radius, w

This is similar to the derivatives with respect to d ,

$$\frac{\partial A}{\partial r_L} = \frac{\partial A}{\partial w} \quad \frac{\partial A}{\partial \phi_L} = R \frac{\partial A}{\partial w}, \quad (\text{A.3.6})$$

where,

$$\frac{\partial A}{\partial w} = \left(\frac{\gamma}{\nu w} \right)^2 \left[\left| \frac{\partial S_H}{\partial w} \right|_{\rho_T} - \left| \frac{\partial S_H}{\partial w} \right|_{\rho_o} \right] - \frac{2A(R, \epsilon)}{w}, \quad (\text{A.3.7})$$

and,

$$\frac{\partial S_H}{\partial w} = \begin{cases} 0 & |w - d| \geq a + \rho, \\ 2\pi w \min[a^2, \rho^2] & w + d \leq |a - \rho|, \\ 2w\mathcal{A}(a, \rho, w + d) & \text{otherwise.} \\ + \frac{2}{\pi} \int_{\max[|w-d|, |a-\rho|]}^{\min[w+d, a+\rho]} \frac{w}{\mu} \cos^{-1} \left(\frac{w^2 + d^2 - \mu^2}{2wd} \right) \sqrt{\Upsilon(a, \rho, \mu)} d\mu & \end{cases} \quad (\text{A.3.8})$$

A.3.3 Image size, ρ

$$\frac{\partial A}{\partial r_T} = \frac{\gamma}{\nu w^2} \left| \frac{\partial S_H}{\partial \rho} \right|_{\rho_T} \quad \frac{\partial A}{\partial r_o} = -\frac{\gamma}{\nu w^2} \left| \frac{\partial S_H}{\partial \rho} \right|_{\rho_o}. \quad (\text{A.3.9})$$

Differentiating (A.7.18)

$$\begin{aligned} \frac{\partial S_H}{\partial \rho} = & \frac{4\rho}{\pi} \int_{\max[|w-d|, |a-\rho|]}^{\min[w+d, a+\rho]} \cos^{-1} \left(\frac{\mu^2 + \rho^2 - a^2}{2\mu\rho} \right) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu \\ & + 4\rho \int_{|w-d|}^{a-\rho} \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu \\ & + 4\rho \int_{|a-\rho|}^{w-d} \cos^{-1} \left(\frac{\mu^2 + \rho^2 - a^2}{2\mu\rho} \right) \mu d\mu + 4\pi\rho \int_0^{\min[w-d, a-\rho]} \mu d\mu. \end{aligned} \quad (\text{A.3.10})$$

Hence,

$$\frac{\partial S_H}{\partial \rho} = \begin{cases} 0 & d - w \geq a + \rho, \\ 2\pi w^2 \rho & w + d \leq a - \rho, \\ 2\pi a^2 \rho & w - d \geq a + \rho, \\ \frac{4\rho}{\pi} \int_{a-\rho}^{\min[w+d, a+\rho]} \cos^{-1} \left(\frac{\mu^2 + \rho^2 - a^2}{2\mu\rho} \right) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu & |w - d| \leq a - \rho < w + d, \\ + 4\rho \mathbb{C}(a - \rho, w, d) & \\ \frac{4\rho}{\pi} \int_{w-d}^{\min[w+d, a+\rho]} \cos^{-1} \left(\frac{\mu^2 + \rho^2 - a^2}{2\mu\rho} \right) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu & |a - \rho| < w - d < a + \rho, \\ + 4\rho \mathbb{C}(w - d, a, \rho) & \\ \frac{4\rho}{\pi} \int_{\max[|w-d|, |a-\rho|]}^{\min[w+d, a+\rho]} \cos^{-1} \left(\frac{\mu^2 + \rho^2 - a^2}{2\mu\rho} \right) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu & \text{otherwise.} \end{cases} \quad (\text{A.3.11})$$

A.3.4 Field of view, a

$$\frac{\partial A}{\partial r_a} = \frac{\gamma R}{f(\nu w)^2} \left[\left| \frac{\partial S_H}{\partial a} \right|_{\rho_T} - \left| \frac{\partial S_H}{\partial a} \right|_{\rho_o} \right]. \quad (\text{A.3.12})$$

This is equivalent to $\frac{\partial S_H}{\partial \rho}$,

$$\frac{\partial S_H}{\partial a} = \begin{cases} 0 & d - w \geq a + \rho, \\ 2\pi w^2 a & w + d \leq \rho - a, \\ 2\pi \rho^2 a & w - d \geq a + \rho, \\ \frac{4a}{\pi} \int_{\rho-a}^{\min[w+d, a+\rho]} \cos^{-1}\left(\frac{\mu^2+a^2-\rho^2}{2\mu a}\right) \cos^{-1}\left(\frac{\mu^2+d^2-w^2}{2\mu d}\right) \mu d\mu & |w-d| \leq \rho-a < w+d, \\ + 4a \mathbb{C}(\rho-a, w, d) & \\ \frac{4a}{\pi} \int_{w-d}^{\min[w+d, a+\rho]} \cos^{-1}\left(\frac{\mu^2+a^2-\rho^2}{2\mu a}\right) \cos^{-1}\left(\frac{\mu^2+d^2-w^2}{2\mu d}\right) \mu d\mu & |a-\rho| < w-d < a+\rho, \\ + 4a \mathbb{C}(w-d, \rho, a) & \\ \frac{4a}{\pi} \int_{\max[|w-d|, |a-\rho|]}^{\min[w+d, a+\rho]} \cos^{-1}\left(\frac{\mu^2+a^2-\rho^2}{2\mu a}\right) \cos^{-1}\left(\frac{\mu^2+d^2-w^2}{2\mu d}\right) \mu d\mu & \text{otherwise.} \end{cases} \quad (\text{A.3.13})$$

A.3.5 Telescope focusing, γ and ν

These parameters enter A through γ and ν and so have a rather more complex derivative,

$$\begin{aligned} \frac{\partial A}{\partial f} = 2A(R, \epsilon) \left[\frac{1}{\gamma} \frac{\partial \gamma}{\partial f} - \frac{1}{\nu} \frac{\partial \nu}{\partial f} \right] + \left(\frac{\gamma}{\nu w} \right)^2 \left[\frac{\partial a}{\partial f} \left(\left| \frac{\partial S_H}{\partial a} \right|_{\rho_T} - \left| \frac{\partial S_H}{\partial a} \right|_{\rho_o} \right) \right. \\ \left. + \frac{\partial \rho_T}{\partial f} \left| \frac{\partial S_H}{\partial \rho} \right|_{\rho_T} - \frac{\partial \rho_o}{\partial f} \left| \frac{\partial S_H}{\partial \rho} \right|_{\rho_o} \right], \quad (\text{A.3.14}) \end{aligned}$$

and identically for Δ where,

$$\frac{\partial \gamma}{\partial f} = -\frac{\Delta}{f^2} \quad \frac{\partial \Delta}{\partial f} = \frac{1}{f} \quad (\text{A.3.15})$$

$$\frac{\partial \nu}{\partial f} = \zeta \frac{(R-f)\Delta}{f^3} \quad \frac{\partial \nu}{\partial \Delta} = \zeta \frac{f-R}{f^2} \quad (\text{A.3.16})$$

$$\frac{\partial a}{\partial f} = -\frac{a}{f+\Delta} \quad \frac{\partial a}{\partial \Delta} = -\frac{a}{f+\Delta} \quad (\text{A.3.17})$$

$$\frac{\partial \rho_{T,o}}{\partial f} = \zeta \Delta \frac{R r_{T,o}}{\gamma^2 f^3} (1 + \gamma) \quad \frac{\partial \rho_{T,o}}{\partial \Delta} = -\zeta \frac{R r_{T,o}}{f^2 \gamma^2}, \quad (\text{A.3.18})$$

with $\zeta = \text{sign}(\gamma - \frac{\Delta}{f^2} R)$.

A.3.6 Laser pulse duration, τ_L

$$\frac{\partial E_i^{\text{ra}}}{\partial \tau_L} = \frac{c \tau_b}{2 \tau_L} \frac{\zeta^{\text{ra}} A(R_i, \epsilon) N_i \exp(-\bar{\chi}_i)}{(R_i + \frac{c}{2} \tau_b - \frac{c}{2} \tau_L)(R_i - \frac{c}{2} \tau_L)} - \frac{E_i^{\text{ra}} - E_B^{\text{ra}}}{\tau_L}, \quad (\text{A.3.19})$$

where $\bar{\chi}_i = (\sigma_R^{\text{el}} + \sigma_R^{\text{ra}}) \mathcal{N}_i + \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}}\right) \chi_i$.

A.3.7 Optical efficiency, ζ

$$\frac{\partial E_i^{\text{ra}}}{\partial \zeta^{\text{ra}}} = \frac{E_i^{\text{ra}} - E_B^{\text{ra}}}{\zeta^{\text{ra}}}. \quad (\text{A.3.20})$$

A.3.8 Optical thickness, χ

For these state vector elements,

$$\frac{\partial E_i^{\text{ra}}}{\partial x} = -(E_i^{\text{ra}} - E_B^{\text{ra}}) \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \frac{\partial \chi_i}{\partial x}. \quad (\text{A.3.21})$$

A.3.8.1 Exponential decay configuration

$$\frac{\partial \chi_i}{\partial z_{BL}} = \begin{cases} -\frac{\chi_i}{z_{BL} + H_{FT}} & R_i \leq z_{BL}, \\ \frac{\chi_\infty}{z_{BL} + H_{FT}} [1 - \exp z_H] - \frac{\chi_i}{z_{BL} + H_{FT}} & R_i > z_{BL}, \end{cases} \quad (\text{A.3.22})$$

$$\frac{\partial \chi_i}{\partial H_{FT}} = \begin{cases} -\frac{\chi_i}{z_{BL} + H_{FT}} & R_i \leq z_{BL}, \\ \frac{\chi_\infty}{z_{BL} + H_{FT}} [1 + (z_H - 1) \exp z_H] - \frac{\chi_i}{z_{BL} + H_{FT}} & R_i > z_{BL}, \end{cases} \quad (\text{A.3.23})$$

$$\frac{\partial \chi_i}{\partial \chi_\infty} = \frac{\chi_i}{\chi_\infty}. \quad (\text{A.3.24})$$

A.3.8.2 Error function configuration

$$\frac{\partial \chi_i}{\partial z_{BL}} = \begin{cases} \alpha_- \left[\text{erf} \left(\frac{z_{BL}}{\sigma_{BL}} \right) - \text{erf} z_\sigma \right] & R_i \leq z_{BL}, \\ \alpha_+ (1 - \exp z_H) + \alpha_- \left[\text{erf} \left(\frac{z_{BL}}{\sigma_{BL}} \right) - \text{erf} z_\sigma \exp z_H \right] & R_i > z_{BL}, \end{cases} \quad (\text{A.3.25})$$

$$\frac{\partial \chi_i}{\partial H_{FT}} = \begin{cases} 0 & R_i \leq z_{BL}, \\ \alpha_+ [1 + (z_H - 1) \exp z_H] + \alpha_- \left\{ \mathcal{D} \left(1 - \frac{\sigma_{BL}^2}{2H_{FT}^2} \right) \right. & R_i > z_{BL}, \\ \left. + (z_H - 1) \text{erf} s_\sigma \exp z_H + \frac{\sigma_{BL}}{\sqrt{\pi}} [\exp(z_H - z_\sigma^2) - 1] \right\} & \end{cases} \quad (\text{A.3.26})$$

$$\frac{\partial \chi_i}{\partial \alpha_+} = \begin{cases} R_i & R_i \leq z_{BL}, \\ z_{BL} + H_{FT} (1 - \exp z_H) & R_i > z_{BL}, \end{cases} \quad (\text{A.3.27})$$

$$\frac{\partial \chi_i}{\partial \alpha_-} = \begin{cases} \sigma_{BL} \left[\frac{z_{BL}}{\sigma_{BL}} \text{erf} \left(\frac{z_{BL}}{\sigma_{BL}} \right) - z_\sigma \text{erf} z_\sigma + \frac{\exp \left(-\frac{z_{BL}^2}{\sigma_{BL}^2} \right) - \exp(-z_\sigma^2)}{\sqrt{\pi}} \right] & R_i \leq z_{BL}, \\ H_{FT} (\mathcal{D} - \text{erf} z_\sigma \exp z_H) + z_{BL} \text{erf} \left(\frac{z_{BL}}{\sigma_{BL}} \right) + \frac{\sigma_{BL}}{\sqrt{\pi}} \left[\exp \left(-\frac{z_{BL}^2}{\sigma_{BL}^2} \right) - 1 \right] & R_i > z_{BL}, \end{cases} \quad (\text{A.3.28})$$

$$\frac{\partial \chi_i}{\partial \sigma_{BL}} = \begin{cases} \frac{\alpha_-}{\sqrt{\pi}} \left[\exp\left(-\frac{z_{BL}^2}{\sigma_{BL}^2}\right) - \exp(-z_\sigma^2) \right] & R_i \leq z_{BL}, \\ \alpha_- \left[\frac{\sigma_{BL}}{2H_{FT}} \mathcal{D} + \frac{\exp\left(-\frac{z_{BL}^2}{\sigma_{BL}^2}\right) - \exp(z_H - z_\sigma^2)}{\sqrt{\pi}} \right] & R_i > z_{BL}, \end{cases} \quad (\text{A.3.29})$$

$$\mathcal{D} = \exp\left(\frac{\sigma_{BL}}{H_{FT}}\right)^2 \left[\operatorname{erf}\left(z_\sigma - \frac{\sigma_{BL}}{2H_{FT}}\right) + \operatorname{erf}\left(\frac{\sigma_{BL}}{2H_{FT}}\right) \right]. \quad (\text{A.3.30})$$

The derivatives for χ_∞ are,

$$\frac{\partial \chi_\infty}{\partial z_{BL}} = \alpha_+ + \alpha_- \operatorname{erf}\left(\frac{z_{BL}}{\sigma_{BL}}\right) \quad (\text{A.3.31})$$

$$\frac{\partial \chi_\infty}{\partial H_{FT}} = \alpha_+ + \alpha_- \left\{ \exp\left(\frac{\sigma_{BL}}{2H_{FT}}\right)^2 \left[\operatorname{erf}\left(\frac{\sigma_{BL}}{2H_{FT}}\right) - 1 \right] \left(1 - \frac{\sigma_{BL}^2}{2H_{FT}^2}\right) - \frac{\sigma_{BL}}{H_{FT}\sqrt{\pi}} \right\} \quad (\text{A.3.32})$$

$$\frac{\partial \chi_\infty}{\partial \alpha_+} = z_{BL} + H_{FT} \quad (\text{A.3.33})$$

$$\frac{\partial \chi_\infty}{\partial \alpha_-} = H_{FT} \exp\left(\frac{\sigma_{BL}}{2H_{FT}}\right)^2 \left[\operatorname{erf}\left(\frac{\sigma_{BL}}{2H_{FT}}\right) - 1 \right] + z_{BL} \operatorname{erf}\left(\frac{z_{BL}}{\sigma_{BL}}\right) - \frac{\sigma_{BL}}{\sqrt{\pi}} \quad (\text{A.3.34})$$

$$\frac{\partial \chi_\infty}{\partial \sigma_{BL}} = \alpha_- \left\{ \frac{\sigma_{BL}}{2H_{FT}} \exp\left(\frac{\sigma_{BL}}{2H_{FT}}\right)^2 \left[\operatorname{erf}\left(\frac{\sigma_{BL}}{2H_{FT}}\right) - 1 \right] + \frac{1}{\sqrt{\pi}} \exp\left(-\frac{z_{BL}^2}{\sigma_{BL}^2}\right) \right\}. \quad (\text{A.3.35})$$

A.4 Calibration function parameter errors

The Jacobians with respect to E_B^{ra} , κ , N_i , and χ^m are all as in Section A.2.

A.4.1 Dead time, τ_d

Without comment on forward model errors due to the representation of the dead time correction,

$$\frac{\partial \varphi_i^{\text{ra}}}{\partial \tau_d} = -\frac{\varphi_i^2}{M\tau_b}. \quad (\text{A.4.1})$$

A.4.2 Range, R

$$\frac{\partial \varphi_i^{\text{ra}}}{\partial R_0} = \frac{M}{(1 + \frac{\tau_d}{\tau_b} \varphi_i^{\text{ra}})^2} \frac{\partial E_i^{\text{ra}}}{\partial R} \quad (\text{A.4.2})$$

$$\frac{\partial \varphi_i^{\text{ra}}}{\partial \tau_b} = \frac{\tau_d}{M\tau_b^2} \varphi_i^2 + \frac{M}{(1 + \frac{\tau_d}{\tau_b} \varphi_i^{\text{ra}})^2} \left[\frac{ci}{2} \frac{\partial E_i^{\text{ra}}}{\partial R} + \frac{\frac{c}{2} \zeta^{\text{ra}} N_i A(R_i, \epsilon) \exp(-\bar{\chi}_i)}{(R_i + \frac{c}{2} \tau_b - \frac{c}{2} \tau_L)(R_i + \frac{c}{2} \tau_b)} \right] \quad (\text{A.4.3})$$

$$\frac{\partial E_i^{\text{ra}}}{\partial R} = \frac{E_i^{\text{ra}} - E_B^{\text{ra}}}{A(R_i, \epsilon)} \frac{\partial A}{\partial R} - \alpha_i (E_i^{\text{ra}} - E_B^{\text{ra}}) \left(1 + \frac{\lambda^{\text{el}}}{\lambda^{\text{ra}}} \right) \quad (\text{A.4.4})$$

$$+ \frac{\zeta^{\text{ra}} N_i A(R_i, \varphi) \exp(-\bar{\chi}_i) \frac{\epsilon}{2} \tau_b (\frac{\epsilon}{2} \tau_L - \frac{\epsilon}{2} \tau_b - 2R_i)}{R_i (R_i + \frac{\epsilon}{2} \tau_b - \frac{\epsilon}{2} \tau_L) (R_i + \frac{\epsilon}{2} \tau_b) (R_i - \frac{\epsilon}{2} \tau_L)}$$

$$\frac{\partial A}{\partial R} = 2A(R, \epsilon) \left(\zeta \frac{\Delta}{\nu f^2} - \frac{\phi_L}{w} \right) \quad (\text{A.4.5})$$

$$+ \left(\frac{\gamma}{\nu w} \right)^2 \left[\frac{r_a}{\gamma f} \left(\left| \frac{\partial S_H}{\partial a} \right|_{\rho_T} - \left| \frac{\partial S_H}{\partial a} \right|_{\rho_o} \right) - \zeta \frac{\Delta}{\gamma f^2} \left(r_T \left| \frac{\partial S_H}{\partial \rho} \right|_{\rho_T} - r_o \left| \frac{\partial S_H}{\partial \rho} \right|_{\rho_o} \right) \right.$$

$$\left. + \phi_L \left(\left| \frac{\partial S_H}{\partial w} \right|_{\rho_T} - \left| \frac{\partial S_H}{\partial w} \right|_{\rho_o} \right) + \frac{\phi_{\parallel} (\delta + \phi_{\parallel} R) + \phi_{\perp}^2 R}{d} \left(\left| \frac{\partial S_H}{\partial d} \right|_{\rho_T} - \left| \frac{\partial S_H}{\partial d} \right|_{\rho_o} \right) \right].$$

A.5 Dead time Jacobean matrix

For convenience, define $y^{\text{H}} = E_i + E_B$ and $y^{\text{L}} = \mathcal{F}E_i + E_B$. The derivatives with respect to E_i , \mathcal{F} , and E_B can then be found by simple application of the chain rule.

A.5.1 Nonparalyzable correction

$$\frac{\partial \varphi_i}{\partial y} = \frac{M\tau_b}{(1 + \tau_d y)^2} \quad \frac{\partial \varphi_i}{\partial \tau_d} = -\frac{\varphi_i^2}{M\tau_b} \quad \frac{\partial \varphi_i}{\partial \tau_b} = \frac{\varphi_i}{\tau_b}. \quad (\text{A.5.1})$$

A.5.2 Donovan correction

$$\frac{\partial \varphi_i}{\partial y} = M\tau_b \exp(-\tau_d y) [P_1 + (2P_2 - P_1)\tau_d y + (3P_3 - P_2)\tau_d^2 y^2 - P_3\tau_d^3 y^3] \quad (\text{A.5.2})$$

$$\frac{\partial \varphi_i}{\partial \tau_d} = M\tau_b y \exp(-\tau_d y) [P_2 - P_1 + (2P_2 - P_1)\tau_d y - P_2\tau_d^2 y^2] \quad (\text{A.5.3})$$

$$\frac{\partial \varphi_i}{\partial P_k} = M\tau_b y \exp(-\tau_d y) (\tau_d y)^{k-1} \quad (\text{A.5.4})$$

$$\frac{\partial \varphi_i}{\partial \tau_b} = \frac{\varphi_i}{\tau_b}. \quad (\text{A.5.5})$$

A.6 Simplification of ξ

Considering the central ray of the lens (fig. 4.3), the displacement of the image from the optical axis μ' is,

$$\tan \Phi = \frac{\mu}{R} = \frac{\mu'}{f + \Delta} \quad (\text{A.6.1})$$

$$\therefore \mu' = \frac{(f + \Delta)\mu}{R} \equiv \gamma \frac{f\mu}{R}, \quad (\text{A.6.2})$$

where $\gamma = 1 + \Delta/f$ and Δ is the displacement of the detector beyond the focal plane.

An equivalent expression for the radius of the image in the detector plane can be found by considering similar triangles about the angle ϕ ,

$$\frac{r'_T}{r_T} = \frac{|R' - \Delta|}{R' + f}, \quad (\text{A.6.3})$$

and similarly for r'_o .

By the thin lens formula,

$$\frac{1}{f} = \frac{1}{R} + \frac{1}{R' + f} \quad (\text{A.6.4})$$

$$\therefore R' = \frac{f^2}{R - f}. \quad (\text{A.6.5})$$

Substituting this into (A.6.3),

$$r'_T = r_T \left| \frac{f^2}{R - f} - \Delta \right| \bigg/ \frac{Rf}{R - f} \quad (\text{A.6.6})$$

$$= \nu \frac{fr_T}{R}, \quad (\text{A.6.7})$$

where $\nu = |\gamma - f^{-2}\Delta R|$.

Using (A.6.2) and (A.6.7) in (4.2.9),

$$\xi(R, \mu, \theta) = \frac{\mathcal{A}(r_a, \nu fr_T/R; \gamma f\mu/R) - \mathcal{A}(r_a, \nu fr_o/R; \gamma f\mu/R)}{\pi (\nu fr_T/R)^2}. \quad (\text{A.6.8})$$

A brief consideration is then given to the form of $\mathcal{A}(r_1, r_2; \mu)$. In the event the circles completely overlap or do not overlap at all, its evaluation is trivially,

$$\mathcal{A}(r_1, r_2; \mu) = \begin{cases} 0 & \mu \geq r_1 + r_2, \\ \pi \min[r_1^2, r_2^2] & \mu \leq |r_1 - r_2|. \end{cases} \quad (\text{A.6.9})$$

Otherwise, consider figure A.1. The shaded area A_1 is obtained by subtracting the area of the dashed triangle from that of the sector,

$$A_1 = r_1^2 \cos^{-1} \left(\frac{x^*}{r_1} \right) - x^* y^*. \quad (\text{A.6.10})$$

Similarly, the area of the other sector is,

$$A_2 = r_2^2 \cos^{-1} \left(\frac{r - x^*}{r_2} \right) - (r - x^*) y^*. \quad (\text{A.6.11})$$




Figure A.1: The area of overlap of two circles, radii r_1 and r_2 , with a distance μ between their centres.

Adding these,

$$\mathcal{A}(r_1, r_2; \mu) = r_1^2 \cos^{-1} \left(\frac{x^*}{r_1} \right) + r_2^2 \cos^{-1} \left(\frac{\mu - x^*}{r_2} \right) - r y^*. \quad (\text{A.6.12})$$

The values of x^* and y^* are derived by solving,

$$r_1^2 = x^{*2} + y^{*2} \quad r_2^2 = (x^* + \mu)^2 + y^{*2} \quad (\text{A.6.13})$$

$$\therefore \quad x^* = \frac{\mu^2 + r_1^2 - r_2^2}{2\mu} \quad y^* = \frac{\sqrt{[(r_1 + r_2)^2 - \mu^2][\mu^2 - (r_1 - r_2)^2]}}{2\mu}. \quad (\text{A.6.14})$$

Hence,

$$\mathcal{A}(r_1, r_2; \mu) = \begin{cases} 0 & \mu \geq r_1 + r_2, \\ \pi \min[r_1^2, r_2^2] & \mu \leq |r_1 - r_2|, \\ \mathbb{A}(r_1, r_2; \mu) & \text{otherwise,} \end{cases} \quad (\text{A.6.15})$$

$$\mathbb{A}(r_1, r_2; \mu) = r_1^2 \cos^{-1} \left(\frac{\mu^2 + r_1^2 - r_2^2}{2\mu r_1} \right) + r_2^2 \cos^{-1} \left(\frac{\mu^2 + r_2^2 - r_1^2}{2\mu r_2} \right) - \frac{1}{2} \sqrt{\Upsilon(r_1, r_2, \mu)} \quad (\text{A.6.16})$$

$$\Upsilon(r_1, r_2, \mu) = 2(r_1^2 r_2^2 + r_1^2 \mu^2 + r_2^2 \mu^2) - r_1^4 - r_2^4 - \mu^4. \quad (\text{A.6.17})$$

Using the symmetry of this expression,

$$\mathcal{A}(r_1, cr_2; c\mu) \equiv c^2 \mathcal{A}(c^{-1}r_1, r_2; \mu). \quad (\text{A.6.18})$$

Hence, equation (A.6.8) can be rewritten as,

$$\xi(R, \mu, \Theta) = \frac{1}{\pi} \left(\frac{\gamma}{\nu r_T} \right)^2 \left[\mathcal{A} \left(\frac{R r_a}{\gamma f}, \frac{\nu}{\gamma} r_T; \mu \right) - \mathcal{A} \left(\frac{\nu R}{\gamma f} r_a, \frac{\nu}{\gamma} r_o; \mu \right) \right]$$

$$\equiv \frac{1}{\pi} \left(\frac{\gamma}{\nu r_T} \right)^2 [\mathcal{A}(a, \rho_T; \mu) - \mathcal{A}(a, \rho_o; \mu)], \quad (\text{A.6.19})$$

where $a = Rr_a/\gamma f$, the radius of the FOV in the plane of the scatterer, and $\rho_{T,o} = \gamma^{-1} \nu r_{T,o}$, the radius in that plane from which illumination reaches a particular point on the detector.

A.7 Evaluation of S_H

Integrating (4.2.17) by parts,

$$S_H(\rho, a, w, d) = \frac{1}{\pi} \left\{ [\mathcal{A}(a, \rho; \mu) \mathcal{A}(\mu, w; d)]_0^{w+d} - \int_{\mu=0}^{w+d} \mathcal{A}(\mu, w; d) d\mathcal{A}(a, \rho; \mu) \right\}. \quad (\text{A.7.1})$$

It can be shown that,

$$d\mathcal{A}(a, \rho; \mu) = \begin{cases} -\sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu} & |a - \rho| \leq \mu \leq a + \rho, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.7.2})$$

The limits of (A.7.1) are thus altered such that $d\mathcal{A} \neq 0$. Using (A.6.17),

$$S_H(\rho, a, w, d) = w^2 \mathcal{A}(a, \rho; w+d) + \frac{1}{\pi} \int_{|a-\rho|}^{\min[w+d, a+\rho]} \mathcal{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu}. \quad (\text{A.7.3})$$

This corrects Halldórsson and Langerholc [1978] as published. They then evaluated this integral numerically to determine the effective area. That is an expensive calculation and often does not converge due to sharp gradients at the boundaries of \mathcal{A} . The efficiency of this calculation can be significantly improved by evaluating the integral where it is analytic.

Firstly, the integrand is zero where $\mu \leq d - w$. Hence, the lower limit should be changed to,

$$S_H(\rho, a, w, d) = w^2 \mathcal{A}(a, \rho; w+d) + \frac{1}{\pi} \int_{\max[d-w, |a-\rho|]}^{\min[w+d, a+\rho]} \mathcal{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu}. \quad (\text{A.7.4})$$

Then, as an illustration¹, the integral is split by the cases of \mathcal{A} ,

$$S_H(\rho, a, w, d) = w^2 \mathcal{A}(a, \rho; w+d) + \frac{1}{\pi} \int_{\max[|w-d|, |a-\rho|]}^{\min[w+d, a+\rho]} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu}$$

¹Herein, if the lower limit of the integral is larger than the upper limit, it will be considered to not exist. Though not strictly correct notation, it eliminates the immediate need for case structures.

$$+ \int_{\max[d-w, |a-\rho|]}^{\min[w-d, a+\rho]} \mu \sqrt{\Upsilon(a, \rho, \mu)} d\mu. \quad (\text{A.7.5})$$

In this form, it is clear that the last term will not exist if $w \leq d$, $d - w \geq a + \rho$, or $w - d \leq |a - \rho|$. The second term ceases to exist when $|w - d| \geq a + \rho$ or $w + d \leq |a - \rho|$ and the first equals zero when $w + d \geq a + \rho$.

The function can be evaluated by then considering the various interactions of these limits. Sketches of the relative positions of the FOV and beam for each case to follow are shown in figure 4.5. For $d - w \geq a + \rho$, all terms disappear as there is no overlap of the beam and detector,

$$S_H^{(1)}(\rho, a, w, d) = 0. \quad (\text{A.7.6})$$

When $w + d \leq |a - \rho|$, neither integral term exists as the beam is entirely within the detector's FOV and the first term evaluates as,

$$S_H^{(2)}(\rho, a, w, d) = \pi w^2 \min[a^2, \rho^2]. \quad (\text{A.7.7})$$

For $w - d \geq a + \rho$, the entire FOV is illuminated and,

$$S_H^{(3)}(\rho, a, w, d) = \frac{1}{\pi} \int_{|a-\rho|}^{a+\rho} \mu \sqrt{\Upsilon(a, \rho, \mu)} d\mu \quad (\text{A.7.8})$$

$$= \left[\frac{\mu^2 - a^2 - \rho^2}{4} \sqrt{\Upsilon(a, \rho, \mu)} + 2a^2 \rho^2 \tan^{-1} \sqrt{\frac{\mu^2 - (a - \rho)^2}{(a + \rho)^2 - \mu^2}} \right]_{|a-\rho|}^{a+\rho} \quad (\text{A.7.9})$$

$$= \pi a^2 \rho^2. \quad (\text{A.7.10})$$

The remaining cases will contain the second term of (A.7.5), which has no analytic solution. To minimise the error and computational cost of evaluating it numerically, it is preferable that any numerical integration be a small correction within a larger expression. That could be the first term of (A.7.5) which equals $w^2 \mathbb{A}(a, \rho, w - d)$ for $|a - \rho| < w - d < a + \rho$, under which conditions the third term of (A.7.5) reduces to,

$$\mathbb{B} = \frac{1}{\pi} \int_{|a-\rho|}^{w-d} \mu \sqrt{\Upsilon(a, \rho, \mu)} d\mu \quad (\text{A.7.11})$$

$$= \frac{(w - d)^2 - a^2 - \rho^2}{4} \sqrt{\Upsilon(a, \rho, w - d)} + 2a^2 \rho^2 \tan^{-1} \sqrt{\frac{(w - d)^2 - (a - \rho)^2}{(a + \rho)^2 - (w - d)^2}}, \quad (\text{A.7.12})$$

if $w - d > |a - \rho|$ and zero otherwise.

Hence, if $|w - d| \leq |a - \rho|$,

$$S_H^{(4)}(\rho, a, w, d) = w^2 \mathbb{A}(a, \rho; w + d) + \frac{1}{\pi} \int_{|a-\rho|}^{w+d} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu}. \quad (\text{A.7.13})$$

If $d - w > |a - \rho|$,

$$S_H^{(5)}(\rho, a, w, d) = w^2 \mathbb{A}(a, \rho; w + d) + \frac{1}{\pi} \int_{d-w}^{w+d} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu}. \quad (\text{A.7.14})$$

If $w - d > |a - \rho|$,

$$S_H^{(6)}(\rho, a, w, d) = w^2 \mathbb{A}(a, \rho; w + d) + \frac{1}{\pi} \int_{w-d}^{w+d} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu} + \mathbb{B}. \quad (\text{A.7.15})$$

When the first term of (A.7.5) equals zero, the same integral may still be used if \mathbb{B} exists, which is true where $w + d \geq a + \rho$ and $|a - \rho| < w - d < a + \rho$,

$$S_H^{(7)}(\rho, a, w, d) = \mathbb{B} + \frac{1}{\pi} \int_{w-d}^{a+\rho} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu}. \quad (\text{A.7.16})$$

The remaining possible cases would, in this formulation, give a single integral term. That can be avoided by returning to (4.2.17) and noting that,

$$d\mathcal{A}(\mu, w; d) = \begin{cases} 0 & \mu \geq w + d \text{ or } \mu \leq d - w, \\ 2\pi\mu d\mu & \mu \leq w - d, \\ 2\mu \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) d\mu & \text{otherwise.} \end{cases} \quad (\text{A.7.17})$$

Hence,

$$\begin{aligned} S_H(\rho, a, w, d) &= \frac{2}{\pi} \int_{|w-d|}^{w+d} \mathcal{A}(a, \rho; \mu) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu + 2 \int_0^{w-d} \mathcal{A}(a, \rho; \mu) \mu d\mu \\ &= \frac{2}{\pi} \int_{\max[|w-d|, |a-\rho|]}^{\min[w+d, a+\rho]} \mathbb{A}(a, \rho; \mu) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu \\ &\quad + 2 \min[a^2, \rho^2] \int_{|w-d|}^{\min[w+d, |a-\rho|]} \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu \\ &\quad + 2 \int_{|a-\rho|}^{w-d} \mathbb{A}(a, \rho; \mu) \mu d\mu + 2\pi \min[a^2, \rho^2] \int_0^{\min[w-d, |a-\rho|]} \mu d\mu. \end{aligned} \quad (\text{A.7.18})$$

In the case that $w + d \geq a + \rho$ and $|w - d| < |a - \rho|$,

$$S_H^{(8)}(\rho, a, w, d) = \frac{2}{\pi} \int_{|a-\rho|}^{a+\rho} \mathbb{A}(a, \rho; \mu) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu + 2 \min[a^2, \rho^2] \mathbb{C}(|a - \rho|, w, d), \quad (\text{A.7.19})$$

where,

$$\mathbb{C}(x, w, d) = \int_{|w-d|}^x \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu + \pi \int_0^{\min[w-d, x]} \mu d\mu \quad (\text{A.7.20})$$

$$= \left[\frac{\mu^2}{2} \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) - \frac{\sqrt{\Upsilon(\mu, w, d)}}{4} - \frac{w^2}{2} \tan^{-1} \left(\frac{w^2 + d^2 - \mu^2}{\sqrt{\Upsilon(\mu, w, d)}} \right) \right]_{|w-d|}^x + \mathcal{H}(w - d) \frac{\pi(w - d)^2}{2} \quad (\text{A.7.21})$$

$$= \frac{x^2}{2} \cos^{-1} \left(\frac{x^2 + d^2 - w^2}{2xd} \right) - \frac{\sqrt{\Upsilon(x, w, d)}}{4} - \frac{w^2}{2} \tan^{-1} \left(\frac{w^2 + d^2 - x^2}{\sqrt{\Upsilon(x, w, d)}} \right) + \frac{\pi w^2}{4}. \quad (\text{A.7.22})$$

The final case, $w + d \geq a + \rho$ and $|a - \rho| < d - w < a + \rho$, can only be cast as a lone integral,

$$S_H^{(9)}(\rho, a, w, d) = \frac{2}{\pi} \int_{d-w}^{a+\rho} \mathbb{A}(a, \rho; \mu) \cos^{-1} \left(\frac{\mu^2 + d^2 - w^2}{2\mu d} \right) \mu d\mu \quad (\text{A.7.23})$$

$$\equiv \frac{1}{\pi} \int_{d-w}^{a+\rho} \mathbb{A}(\mu, w; d) \sqrt{\Upsilon(a, \rho, \mu)} \frac{d\mu}{\mu}. \quad (\text{A.7.24})$$

Fortunately, this integration appears to be stable (if slow).

Case 8 deviates from Povey et al. [2012], which instead quoted case 9 twice (with different lower limits). The new form addresses some failures of the integration to converge when $d < 0$. Originally, S_H was simply taken to equal zero in such cases. That approximation was not found to have a significant effect on the result, but its removal is clearly preferable as it gives a more consistent model and vastly decreases the chance of an integration failing to converge.

A.8 IDL computer codes

This section presents excerpts of the IDL source code for the forward models and the optimal estimation retrieval scheme. They are presented ‘as is’ and are not intended to be directly applied by the reader, as they reference several local sub-routines. They should merely provide some guidance as to the practical concerns of the techniques.

A.8.1 Scattering forward model

```

PRO MODEL_PROFILE7, x_, num, v0, v1, nex, nr, facs, r_x, r_y, back, fmy, k, Df.Dcoef, o0, o1, kb, log=log_, noise=shots,
    inta_y=inta_y, ratio=ratio_, mod_0=mod_0_, mod_1=mod_1_
log=KEYWORD.SET(log_)
ratio=KEYWORD.SET(ratio_)
if log then begin
    if ratio then x=TRANPOSE([EXP(x_[0:nex-1]),x_[nex :*]]) else x=EXP(x_)
endif else x=x_

; set-up interpolation
MODEL_SPLINE,r_x,r_y,j,int,ef,der

; interpolate backscatter onto the y axis
bpb=[0d0,TRANPOSE(MATRIX.MULTIPLY(x[0:nex-1],ef,/at)),0d0]
beta_y=int[* ,0]*x[j]+int[* ,1]*x[j+1]+int[* ,2]*bpb[j]+int[* ,3]*bpb[j+1]

; integrate extinction over r_x with trapezium rule
inta_x=INT_TRAP(r_x, ratio ? x[0:nex-1]*x[nex :2*nex-1] : x[nex :2*nex-1])
ipp=[0d0,TRANPOSE(MATRIX.MULTIPLY(inta_x,ef,/at)),0d0]
inta_y=int[* ,0]*inta_x[j]+int[* ,1]*inta_x[j+1]+int[* ,2]*ipp[j] $
    +int[* ,3]*ipp[j+1]

; evaluate lidar equation
term0=x[2*nex]*v0*EXP(-2d0*inta_y)
mod_0=(facs[0]*num+beta_y)*term0
mod_1=x[2*nex]*v1*EXP(-facs[1]*inta_y)
mod_0=mod_0+back[0,0]
mod_1=mod_1+back[0,1]

; PMT correction
t0=1d0+facs[2]*mod_0 & t1=1d0+facs[3]*mod_1
fmy=TRANPOSE([mod_0/t0,mod_1/t1])

; add Poisson noise (for use as a signal generator)
if N_ELEMENTS(shots) eq 1 then for i=0,2*nr-1 do $
    fmy[i]=TOTAL(RANDOMN(seed,shots,poisson=fmy[i]/DOUBLE(shots)))

; see p.264
; covariance matrix, d(fmy[row])/d(x[col])
if ARG_PRESENT(k) then begin
    k=DBLARR(2*nex+1,2*nr)
    temp0=DBLARR(nex,nr)
    temp1=DBLARR(nex,nr)

; elastic backscatter
i=0
while j[i] eq 0 do begin ; skip c term as y''_0=0
    temp0[* ,i]=int[i,3]*ef[* ,0]
    temp0[0,i]=int[i,0]
    temp0[1,i]=int[i,1]
    ++i
endwhile
while j[i] lt nex-2 do begin
    temp0[* ,i]=int[i,2]*ef[* ,j[i]-1]+int[i,3]*ef[* ,j[i]]
    temp0[j[i],i]=int[i,0]
    temp0[j[i]+1,i]=int[i,1]
    ++i
endwhile
while i lt nr do begin ; skip d term as y''_{n-1}=0
    temp0[* ,i]=int[i,2]*ef[* ,nex-3]
    temp0[nex-2,i]=int[i,0]
    temp0[nex-1,i]=int[i,1]
    ++i
endwhile

; extinction
p=nex-1-LINDGEN(nex)
q=1+LINDGEN(nex-2)
for i=0,nr-1 do begin
    temp2=(TOTAL(temp0[p,i+LONARR(nex)],/cum,/double))[p]

    temp1[nex-1,i]=5d-1*(r_x[nex-1]-r_x[nex-2])*temp0[nex-1,i]
    temp1[q,i]=5d-1*(r_x[q]-r_x[q-1])*temp0[q,i] + 5d-1*(r_x[q+1]-r_x[q-1])*temp2[q+1]
    temp1[0,i]=r_x[0]*temp0[0,i] + 5d-1*(r_x[1]+r_x[0])*temp2[1]
endifor

; evaluation of K
if ratio then begin
    for i=0,nex-1 do begin
        k[i,0:nr-1]=term0*temp0[i,*]-2d0*x[nex+i]*mod_0*temp1[i,*]
        k[i,nr :*]=-facs[1]*x[nex+i]*mod_1*temp1[i,*]
        k[nex+i,0:nr-1]=-2d0*x[i]*mod_0*temp1[i,*]
        k[nex+i,nr :*]=x[i]/x[nex+i]*k[i,nr :*]
    endfor
endif else begin
    for i=0,nex-1 do begin
        k[i,0:nr-1]=term0*temp0[i,*]
        k[nex+i,0:nr-1]=-2d0*mod_0*temp1[i,*]
        k[nex+i,nr :*]=-facs[1]*mod_1*temp1[i,*]
    endfor
endif

```

```

endelse

; power
k[2*nex,0:nr-1]=mod.0./x[2*nex]
k[2*nex,nr :*] =mod.1./x[2*nex]

; PMT correction to K, see p.270 (and p.340 for correction)
for i=0,nr-1 do begin
  k[*,i]/=t0[i]*t0[i]
  k[*,i+nr]/=t1[i]*t1[i]
endfor

; log correction
if log then for i=0, ratio ? nex-1:2*nex do k[i,*]=x[i]
endif

if ARG.PRESENT(kb) then begin
  ; parameter error Jacobean calculation
  l=INDGEN(nr)
  kb=DBLARR(4*nr+7,2*nr)

  kb[0,0:nr-1]=-facs[6]*fmy[0:nr-1]*fmy[0:nr-1] ; tau[0]
  kb[1,nr :*]=-facs[6]*fmy[nr :*]*fmy[nr :*] ; tau[1]
  kb[2,0:nr-1]= 1d0 ; back[0,0]
  kb[3,nr :*] = 1d0 ; back[0,1]

  temp2=der[*,0]*inta_x[j]+der[*,1]*inta_x[j+1]+der[*,2]*ipp[j]+der[*,3]*ipp[j+1]
  kb[4,0:nr-1]=-2d0*mod.0.*(1d0/r.y+temp2) $
    +term0*(der[*,0]*x[j]+der[*,1]*x[j+1]+der[*,2]*bpp[j]+der[*,3]*bpp[j+1]) ; yoff
  kb[4,nr :*]=-mod.1.*(2d0/r.y+facs[1]*temp2) ; yoff
  kb[5,nr :*]=-1.49896229d-1*mod.1.*(2d0/r.y+facs[1]*temp2) ; width

  kb[6,nr :*]=-Df.Dcoef*inta_y*mod.1. ; coef

  kb[7+ 1, 1]=mod.0./o0 ; over
  kb[7+ nr+1,nr+1]=mod.1./o1 ; over.n2
  kb[7+2*nr+1, 1]=facs[0]*term0 & kb[7+2*nr+1,nr+1]=mod.1./num ; num
  kb[7+3*nr+1, 1]=-facs[4]*mod.0. & kb[7+3*nr+1,nr+1]=-facs[5]*mod.1. ; ray

  for i=0,nr-1 do begin ; PMT correction
    kb[2:*,i]/=t0[i]*t0[i]
    kb[2:*,i+nr]/=t1[i]*t1[i]
  endfor
endif
END

```

A.8.2 Calibration function forward model

```

PRO CALIB, x, inst., r, nr, num, ray, back, facs, log, f, k, kb, exp=exp, cal=cal, over=over, chi=chi, grad=g.cal, kk=
  kk
inst=inst.
inst.delta=log[0] ? EXP(x[0]):x[0]
inst.phi1=log[1] ? EXP(x[1]):x[1]
inst.phi2=log[2] ? EXP(x[2]):x[2]
inst.ldia=log[3] ? EXP(x[3]):x[3]
inst.ldiv=log[4] ? EXP(x[4]):x[4]
inst.tdia=log[5] ? EXP(x[5]):x[5]
inst.odia=log[6] ? EXP(x[6]):x[6]
inst.adia=log[7] ? EXP(x[7]):x[7]
inst.flength=log[8] ? EXP(x[8]):x[8]
inst.dis=log[9] ? EXP(x[9]):x[9]
inst.ldur=log[10] ? EXP(x[10]):x[10]
eta=log[11] ? EXP(x[11]):x[11]
z.bl=log[12] ? EXP(x[12]):x[12]
h.ft=log[13] ? EXP(x[13]):x[13]
alp0=log[14] ? EXP(x[14]):x[14] ; = AOT(\infty) for \EXP mode, \alpha+ otherwise

; calibration function
c=1.49896229d-1
r.b=c*inst.width
r.l=c*inst.ldur
cal=inst.nshots*eta/r.l*ALOG(r*(r+r.b-r.l)/((r+r.b)*(r-r.l)))
over = ARG.PRESENT(k) ? AREA4(inst,r,grad=g,k=kk) : AREA4(inst,r,k=kk)

; aerosol optical thickness
chi=DBLARR(nr)
z.h=(z.bl-r)/h.ft
p=WHERE(r.le z.bl,np,comp=q,ncomp=nq)
if KEYWORD.SET(exp) then begin
  ext=alp0/(z.bl+h.ft)
  if np.gt.0 then chi[p]=ext*r[p]
  if nq.gt.0 then chi[q]=ext*(z.bl+h.ft*(1d0-EXP(z.h[q])))
  aot=alp0
endif else begin
  alp1=log[15] ? EXP(x[15]):x[15] ; not used for \EXP mode
  s.bl=log[16] ? EXP(x[16]):x[16] ; not used for \EXP mode
  z.s=(z.bl-r)/s.bl
  z.temp=z.bl/s.bl
  s.temp=5d-1*s.bl/h.ft
  if np.gt.0 then begin

```

```

term2=(EXP(-z.temp*z.temp)-EXP(-z.s[p]*z.s[p]))/SQRT(!dpi)
chi[p]=alp0*r[p] + alp1*s.bl*(z.temp*ERF(z.temp)-z.s[p]*ERF(z.s[p])+term2)
endif
if nq gt 0 then begin
  test=EXP(s.temp*s.temp)
  if ~FINITE(1d3*test) then test=0d0
  term1=test*(ERF(z.s[q]-s.temp)+ERF(s.temp))
  chi[q]=alp0*(z.bl+h.ft*(1d0-EXP(z.h[q]))) $
    +alp1*(h.ft*term1 - h.ft*EXP(z.h[q])*ERF(z.s[q]) + z.bl*ERF(z.temp) $
    + s.bl/SQRT(!dpi)*(EXP(-z.temp*z.temp)-1d0))
endif
aot=alp0*(z.bl+h.ft) + alp1*(h.ft*test*(ERF(s.temp)-1d0)+z.bl*ERF(z.temp) $
+ s.bl/SQRT(!dpi)*(EXP(-z.temp*z.temp)-1d0))
endelse

; Raman lidar equation
short=num*ray*EXP(-facs[0]*chi)
mod.l.=cal*over*short
mod.l.=mod.l.+back
t1=1d0+facs[1]*mod.l & t12=1d0/(t1*t1)
f=TRANPOSE([mod.l/t1,aot])

if ARG_PRESENT(k) then begin
  k=DBLARR(exp ? 15:17,nr+1)
  kb=DBLARR(5+2*nr,nr+1)
  g_cal=DBLARR(12,nr)
  k[0:9,0:nr-1]=MM(g[0:9,*],cal*num*ray*EXP(-facs[0]*chi),/r)
  k[10,0:nr-1]=inst.nshots*eta*r.b/inst.ldur/((r+r.b-r.l)*(r-r.l))*over*short - mod.l./inst.ldur
  k[11,0:nr-1]=mod.l./eta

  g_cal[0:9,*]=MM(g[0:9,*],cal,/r)
  g_cal[10,*]=(inst.nshots*eta*r.b/inst.ldur/((r+r.b-r.l)*(r-r.l)) - cal/inst.ldur)*over
  g_cal[11,*]=over*cal/eta

  if KEYWORD.SET(exp) then begin
    if np gt 0 then begin
      k[12,p]=chi[p]/(z.bl+h.ft)
      k[13,p]=k[12,p]
      kb[2,p]=-facs[0]*ext*mod.l.[p]*t12[p]
    endif
    if nq gt 0 then begin
      k[12,q]=ext*(1d0-EXP(z.h[q]))-chi[q]/(z.bl+h.ft)
      k[13,q]=ext*(1d0+(z.h[q]-1d0)*EXP(z.h[q]))-chi[q]/(z.bl+h.ft)
      kb[2,q]=-facs[0]*ext*EXP(z.h[q])*mod.l.[q]*t12[q]
    endif
    k[14,0:nr-1]=chi/alp0
    k[12:14,0:nr-1]=MM(k[12:14,0:nr-1],-facs[0]*mod.l.,/r)

    k[14,nr]=1d0 ; AOT derivative
  endif else begin
    if np gt 0 then begin
      k[12,p]=alp1*(ERF(z.temp)-ERF(z.s[p]))
      k[14,p]=r[p]
      k[15,p]=s.bl*(z.temp*ERF(z.temp) - z.s[p]*ERF(z.s[p]) + term2)
      k[16,p]=alp1*term2
      kb[2,p]=-facs[0]*mod.l.[p]*(alp0+alp1*ERF(z.s[p]))*t12[p]
    endif
    if nq gt 0 then begin
      k[12,q]=alp0*(1d0-EXP(z.h[q]))+alp1*(ERF(z.temp)-ERF(z.s[q])*EXP(z.h[q]))
      k[13,q]=alp0*(1d0+(z.h[q]-1d0)*EXP(z.h[q])) $
        + alp1*(term1*(1d0-2d0*s.temp*s.temp) + (z.h[q]-1)*EXP(z.h[q])*ERF(z.s[q]) $
        + s.bl/(SQRT(!dpi)*h.ft)*(EXP(z.h[q]-z.s[q]*z.s[q])-1))
      k[14,q]=z.bl+h.ft*(1d0-EXP(z.h[q]))
      k[15,q]=h.ft*term1 - h.ft*EXP(z.h[q])*ERF(z.s[q]) + z.bl*ERF(z.temp) $
        + s.bl/SQRT(!dpi)*(EXP(-z.temp*z.temp)-1d0)
      k[16,q]=alp1*(s.temp*term1 + (EXP(-z.temp*z.temp)-EXP(z.h[q]-z.s[q]*z.s[q]))/SQRT(!dpi))
      kb[2,q]=-facs[0]*mod.l.[q]*EXP(z.h[q])*(alp0+alp1*ERF(z.s[q]))*t12[q]
    endif
    k[12:16,0:nr-1]=MM(k[12:16,0:nr-1],-facs[0]*mod.l.,/r)

    ; AOT derivatives
    k[12,nr]=alp0+alp1*ERF(z.temp)
    k[13,nr]=alp0+alp1*(test*(ERF(s.temp)-1d0)*(1d0-2d0*s.temp*s.temp) $
    -s.bl/(h.ft*SQRT(!dpi)))
    k[14,nr]=z.bl+h.ft
    k[15,nr]=h.ft*test*(ERF(s.temp)-1d0)+z.bl*ERF(z.temp)-s.bl/SQRT(!dpi)
    k[16,nr]=alp1*(s.temp*test*(ERF(s.temp)-1d0)+EXP(-z.temp*z.temp)/SQRT(!dpi))
  endelse

; PMT correction
for i=0,nr-1 do k[*,i]=t12[i]

; log correction
for i=0,exp ? 14:16 do if log[i] then k[i,*]=EXP(x[i])
for i=0,11 do if log[i] then g_cal[i,*]=EXP(x[i])

; parameter error
l=LINDGEN(nr)
kb[0,0:nr-1]=-facs[2]*f[0:nr-1]*f[0:nr-1] ; tau_d
kb[2,0:nr-1]=short*t12*(cal*g[10,*] $ ; yoffset
+ inst.nshots*eta*r.b*over*(r.l-r.b-2d0*r)/(r*(r+r.b-r.l)*(r+r.b)*(r-r.l)))
kb[1,0:nr-1]=facs[3]*f[0:nr-1]*f[0:nr-1]+ $

```

```

c*(1*kb[2,0:nr-1] + inst.nshots*eta*t12*over*short/((r+r_b-r_l)*(r+r_b))) ; width

kb[3,0:nr-1]=t12 ; back
kb[4,0:nr-1]=-facs[4]*chi*mod.1.*t12 ; coef
kb[5+1,1]=mod.1.*t12/num ; num
kb[5+nr+1,1]=mod.1.*t12/ray ; ray
endif
END

```

A.8.3 Effective area model

```

FUNCTION AREA4.UPS, r1, r2, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK, th, alp, rho, w, d, i, j, alp2, rho2, w2, d2
  r12=r1*r1
  r22=r2*r2
  mu2=mu*mu
  if (r12 lt th) || ((r22 lt th) || (mu2 lt th)) then RETURN, 0d0
  RETURN, SQRT((2d0*(r12*r22+r12*mu2+r22*mu2)-r12*r12-r22*r22-mu2*mu2) > 0d0)
END

FUNCTION AREA4.A, r1, r2, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  if ABS(r1+r2-mu) lt th OR ABS(ABS(r1-r2)-mu) lt th then RETURN, 0d0
  r12=r1*r1
  r22=r2*r2
  mu2=mu*mu
  ang0=((mu2+r12-r22)/(2d0*mu*r1) > (-1d0)) < 1d0
  ang1=((mu2-r12+r22)/(2d0*mu*r2) > (-1d0)) < 1d0
  if r2 lt 2d-4*r1 then begin ; expansion to avoid rounding errors
    t=mu-r1
    RETURN, r22*ACOS(ang1)-t*SQRT(r1*(r22-t*t)/mu)
  endif
  if r1 lt 2d-4*r2 then begin
    t=mu-r2
    RETURN, r22*ACOS(ang0)-t*SQRT(r2*(r12-t*t)/mu)
  endif
  RETURN, r12*ACOS(ang0)+r22*ACOS(ang1) $
    -5d-1*SQRT((2d0*(r12*r22+r12*mu2+r22*mu2)-r12*r12-r22*r22-mu2*mu2) > 0d0)
END

FUNCTION AREA4.C, r1, r2, r3 ; 2*(3.5.53) of thesis
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  r12=r1*r1
  r22=r2*r2
  r32=r3*r3
  RETURN, r12*ACOS((r12+r32-r22)/(2d0*r1*r3)) - 5d-1*AREA4.UPS(r1,r2,r3) $
    -r22*ATAN((r32+r22-r12)/AREA4.UPS(r1,r2,r3)) + 5d-1*dpi*r22
END

FUNCTION AREA4.INT1, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  RETURN, AREA4.A(mu,w[i],d[i])*AREA4.UPS(alp[i],rho[i,j],mu)/mu
END

FUNCTION AREA4.INT2, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  RETURN, AREA4.A(alp[i],rho[i,j],mu)*ACOS(((mu*mu+d2[i]-w2[i])/(2d0*mu*d[i]) > (-1d0)) < 1d0)*mu
END

FUNCTION AREA4.DER1, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  RETURN, AREA4.UPS(mu,w[i],d[i])*AREA4.UPS(alp[i],rho[i,j],mu)/mu
END

FUNCTION AREA4.DER2, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  RETURN, ACOS(((w2[i]+d2[i]-mu*mu)/(2d0*w[i]*d[i]) > (-1d0)) < 1d0)*AREA4.UPS(alp[i],rho[i,j],mu)/mu
END

FUNCTION AREA4.DER3, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  mu2=mu*mu
  RETURN, ACOS(((mu2+rho2[i,j]-alp2[i])/(2d0*mu*rho[i,j]) > (-1d0)) < 1d0) $
    *ACOS(((mu2+d2[i]-w2[i])/(2d0*mu*d[i]) > (-1d0)) < 1d0)*mu
END

FUNCTION AREA4.DER4, mu
  COMPILE.OPT HIDDEN
  COMMON A4.BLOCK
  mu2=mu*mu
  RETURN, ACOS(((mu2+alp2[i]-rho2[i,j])/(2d0*mu*alp[i]) > (-1d0)) < 1d0) $
    *ACOS(((mu2+d2[i]-w2[i])/(2d0*mu*d[i]) > (-1d0)) < 1d0)*mu
END

```

```

END

FUNCTION AREA4.GAUSS, x
COMMON A4.BLOCK
  if x le 0d0 OR x ge (alp[i]+rho[i,j]) then RETURN, 0d0

  temp=2d0*x*EXP(-(x*x+d2[i])/w2[i])*BESELI(2d0*x*d[i]/w2[i],0)
  if x le ABS(alp[i]-rho[i,j]) then begin
    if alp[i] le rho[i,j] then RETURN, !dpi*alp2[i]*temp else RETURN, !dpi*rho2[i,j]*temp
  endif else RETURN, AREA4.A(alp[i],rho[i,j],x)*temp
END

FUNCTION AREA4, inst_, r, jmax=jm, eps=eps, k=kk, gradients=grad, threshold=th_, ans=ans, tdia=r.t, odia=r.o, adia=r.a,
  f=f1, dis=dis, delta=delta, ldia=r.l, lddiv=phi.l, phil=phi.para, phi2=phi.perp, gaussian=gaus, limit=lim
COMMON A4.BLOCK
  if ~KEYWORD.SET(eps) then eps=1d-8
  if ~KEYWORD.SET(kk) then kk=9
  if ~KEYWORD.SET(th_) then th=1d-14 else th=th_
  if ~KEYWORD.SET(lim) then lim=1d-13
  n=N.ELEMENTS(r)
  ans=DBLARR(n,2)

  ; manage parameters
  inst=inst_
  if KEYWORD.SET(r.t) then inst.tdia = r.t
  if KEYWORD.SET(r.o) then inst.odia = r.o
  if KEYWORD.SET(r.a) then inst.adia = r.a
  if KEYWORD.SET(f1) then inst.flength=f1
  if KEYWORD.SET(dis) then inst.dis =dis
  if KEYWORD.SET(delta) then inst.delta =delta
  if KEYWORD.SET(r.l) then inst.ldia =r.l
  if KEYWORD.SET(phi.l) then inst.ldiv =phi.l
  if KEYWORD.SET(phi.para) then inst.phil =phi.para
  if KEYWORD.SET(phi.perp) then inst.phi2 =phi.perp

  ; arguments
  f=1d0/inst.flength
  f2=f*f
  gam=1d0+inst.dis*f
  nu=ABS(gam-inst.dis*r*f2)
  alp=5d-1*inst.adia*r*f/gam
  rho=[[5d-1*inst.tdia*nu/gam],[5d-1*inst.odia*nu/gam]]
  w=5d-1*inst.ldia+inst.ldiv*r
  d=SQRT((inst.delta+inst.phil*r)*(inst.delta+inst.phil*r) + $
    inst.phi2*inst.phi2*r*r)
  if gam eq 0 then begin
    grad=DBLARR(11,n)
    RETURN, DBLARR(n)
  endif

  ; convenient shorthands
  x=w+d
  y=w-d
  a=[[alp+rho[* ,0]], [alp+rho[* ,1]]]
  b=[[ABS(alp-rho[* ,0])], [ABS(alp-rho[* ,1])]]
  alp2=alp*alp
  rho2=rho*rho
  w2=w*w
  d2=d*d
  x2=x*x
  y2=y*y
  a2=a*a
  b2=b*b

  ; use Gaussian beam profile without all the clever maths (ignores GRAD)
  if KEYWORD.SET(gaus) then begin
    for j=0,1 do for i=0,n-1 do ans[i,j]=QROMB('AREA4.GAUSS',0,a[i,j],e=eps,j=jm,k=kk,/d)
    temp=gam/(nu*w) & temp2=temp*temp
    RETURN, temp2*(ans[* ,0]-ans[* ,1])
  endif

  ; see thesis (4.5.50) for cases of S.H
  for j=0,1 do begin
    for i=0,n-1 do begin
      if (-y[i]) lt a[i,j] then begin ; S.H case 1 (do nothing)
        if x[i] le b[i,j] then begin ; S.H case 2
          if alp[i] lt rho[i,j] $
            then ans[i,j]=!dpi*w2[i]*alp2[i] $
            else ans[i,j]=!dpi*w2[i]*rho2[i,j]
          endif else begin
            if y[i] ge a[i,j] then begin ; S.H case 3
              ans[i,j]=!dpi*alp2[i]*rho2[i,j]
            endif else begin
              if x[i] ge a[i,j] then begin ; S.H cases 7, 8, 9
                if ABS(y[i]) le b[i,j] then begin ; S.H case 8
                  if alp[i] lt rho[i,j] $
                    then ans[i,j]=alp2[i] $
                    else ans[i,j]=rho2[i,j]
                  ans[i,j]*=AREA4.C(b[i,j],w[i],d[i])
                  ans[i,j]=2d0/!dpi*QROMB('AREA4.INT2',b[i,j],a[i,j],e=eps,j=jm,k=kk,/d)
                endif else begin
                  if w[i] ge d[i] then begin ; S.H case 7

```

```

ans[i,j]=2.5d-1*(y2[i]-alp2[i]-rho2[i,j])*AREA4.UPS(alp[i],rho[i,j],y[i]) $
+2d0*alp2[i]*rho2[i,j]*ATAN(SQRT((y2[i]-b2[i,j])/(a2[i,j]-y2[i]))) ; \blocB
ans[i,j]=QROMB('AREA4.INT1',y[i],a[i,j],e=eps,j=jm,k=kk,/d)/!dpi
endif else begin ; S.H case 9
ans[i,j]=QROMB('AREA4.INT1',-y[i],a[i,j],e=eps,j=jm,k=kk,/d)/!dpi
endifelse
endifelse
endif else begin ; S.H cases 4, 5, 6
ans[i,j]=w2[i]*AREA4.A(alp[i],rho[i,j],x[i]) ; w^2 \blocA
if ABS(y[i]) le b[i,j] then begin ; S.H case 4
ans[i,j]=QROMB('AREA4.INT1',b[i,j],x[i],e=eps,j=jm,k=kk,/d)/!dpi
endif else begin
if w[i] ge d[i] then begin ; S.H case 6
ans[i,j]=2.5d-1*(y2[i]-alp2[i]-rho2[i,j])*AREA4.UPS(alp[i],rho[i,j],y[i]) $
+2d0*alp2[i]*rho2[i,j]*ATAN(SQRT((y2[i]-b2[i,j])/(a2[i,j]-y2[i]))) ; \blocB
ans[i,j]=QROMB('AREA4.INT1',y[i],x[i],e=eps,j=jm,k=kk,/d)/!dpi
endif else begin ; S.H case 5
ans[i,j]=QROMB('AREA4.INT1',-y[i],x[i],e=eps,j=jm,k=kk,/d)/!dpi
endifelse
endifelse
endifelse
endif
endif
endif
endif
endif
endif
temp=gam/(nu*w) & temp2=temp*temp
eff=temp2*(ans[* ,0]-ans[* ,1]) ; effective area (i.e. the answer)
trash=WHERE( FINITE(eff),nf)
if nf gt 0 then MESSAGE, 'Bad effective area.'

; see thesis Appendix for derivatives of S.H
if ARG.PRESENT(grad) then begin
derd=DBLARR(n,2)
derw=DBLARR(n,2)
derp=DBLARR(n,2)
dera=DBLARR(n,2)
for j=0,1 do begin
l=MIN([x],[a[* ,j]]) ,dim=2)
for i=0,n-1 do begin
if (-y[i]) lt a[i,j] then begin ; case 1
if x[i] le b[i,j] then begin ; case 2
if alp[i] lt rho[i,j] then begin
dera[i,j]=2d0*!dpi*w2[i]*alp[i]
derw[i,j]=2d0*!dpi*w[i]*alp2[i]
endif else begin
derp[i,j]=2d0*!dpi*w2[i]*rho[i,j]
derw[i,j]=2d0*!dpi*w[i]*rho2[i,j]
endifelse
endif else begin
if y[i] ge a[i,j] then begin ; case 3
dera[i,j]=2d0*!dpi*alp[i]*rho2[i,j]
derp[i,j]=2d0*!dpi*alp2[i]*rho[i,j]
endif else begin
t=MAX([ABS(y[i]),b[i,j]])
derd[i,j]=QROMB('AREA4.DER1',t,l[i],e=eps,j=jm,k=kk,/d)/(!dpi*d[i])
if x[i] lt a[i,j] then derw[i,j]=2d0*w[i]*AREA4.A(alp[i],rho[i,j],x[i])
derw[i,j]=QROMB('AREA4.DER2',t,l[i],e=eps,j=jm,k=kk,/d)*2d0*w[i]/!dpi

if ABS(y[i]) le b[i,j] then begin
if alp[i] lt rho[i,j] then begin ; a case 4, rho case 6
dera[i,j]=2d0*alp[i]*AREA4.C(b[i,j],w[i],d[i])
if rho2[i,j] gt lim then $
dera[i,j]=QROMB('AREA4.DER4',b[i,j],l[i],e=eps,j=jm,k=kk,/d)*(4d0*alp[i]/!dpi)
if alp2[i] gt lim then $
derp[i,j]=QROMB('AREA4.DER3',b[i,j],l[i],e=eps,j=jm,k=kk,/d)*(4d0*rho[i,j]/!dpi)
endif else begin ; a case 6, rho case 4
derp[i,j]=2d0*rho[i,j]*AREA4.C(b[i,j],w[i],d[i])
if rho2[i,j] gt lim then $
dera[i,j]=QROMB('AREA4.DER4',b[i,j],l[i],e=eps,j=jm,k=kk,/d)*(4d0*alp[i]/!dpi)
if alp2[i] gt lim then $
derp[i,j]=QROMB('AREA4.DER3',b[i,j],l[i],e=eps,j=jm,k=kk,/d)*(4d0*rho[i,j]/!dpi)
endifelse
endif else begin
if w[i] ge d[i] then begin ; case 5
dera[i,j]=2d0*alp[i]*AREA4.C(y[i],rho[i,j],alp[i])
derp[i,j]=2d0*rho[i,j]*AREA4.C(y[i],alp[i],rho[i,j])
if rho2[i,j] gt lim then $
dera[i,j]=QROMB('AREA4.DER4',y[i],l[i],e=eps,j=jm,k=kk,/d)*(4d0*alp[i]/!dpi)
if alp2[i] gt lim then $
derp[i,j]=QROMB('AREA4.DER3',y[i],l[i],e=eps,j=jm,k=kk,/d)*(4d0*rho[i,j]/!dpi)
endif else begin ; case 6
if rho2[i,j] gt lim then $
dera[i,j]=QROMB('AREA4.DER4',-y[i],l[i],e=eps,j=jm,k=kk,/d)*(4d0*alp[i]/!dpi)
if alp2[i] gt lim then $
derp[i,j]=QROMB('AREA4.DER3',-y[i],l[i],e=eps,j=jm,k=kk,/d)*(4d0*rho[i,j]/!dpi)
endifelse
endifelse
endifelse
endif
endif
endif
endif
endif

```

```

endfor

; construct K
grad=DEBLARR(n,11) ; transpose of what it should be for memory efficiency
Ad=temp2*(derd[* ,0]-derd[* ,1])/d ; (A.1.30)/d
grad[* ,0]=Ad*(inst .delta+inst .phi1*r) ; \delta
grad[* ,1]=grad[* ,0]*r ; \phi-\parallel
grad[* ,2]=inst .phi2*r*r*Ad ; \phi-\perp

Aw=temp2*(derw[* ,0]-derw[* ,1])-2d0*eff/w
grad[* ,3]=5d-1*Aw ; r.l (5d-1 represents difference between diameter and radius)
grad[* ,4]=r*Aw ; \phi.L

grad[* ,5]=5d-1*temp/w*derp[* ,0] ; r.T
grad[* ,6]=-5d-1*temp/w*derp[* ,1] ; r.O

s=(gam-inst .dis*r*f2)/nu ; sign(\nu)
grad[* ,7]=5d-1*f/gam*temp2*r*(dera[* ,0]-dera[* ,1]) ; r.a
grad[* ,8]=-2d0*s*(1d0+gam)*inst .dis*f*f2/gam*eff*r/nu $
+temp2*(5d-1*s*inst .dis*f*f2*(1d0+gam)/(gam*gam)*r*(inst .tdia*derp[* ,0]-inst .odia*derp[* ,1]) $
-alp/(inst .flength+inst .dis)*(dera[* ,0]-dera[* ,1])) ; f
grad[* ,9]=2d0*eff*(f/gam-s*f*(1d0-r*f)/nu) $
-temp2*(alp/(inst .flength+inst .dis)*(dera[* ,0]-dera[* ,1]) $
+5d-1*s*f2/(gam*gam)*r*(inst .tdia*derp[* ,0]-inst .odia*derp[* ,1])) ; dis (\Delta)
grad[* ,10]=2d0*eff*(s*inst .dis*f2/nu-inst .ldiv/w) $ ; R.O
+temp2*(5d-1*inst .adia*f/gam*(dera[* ,0]-dera[* ,1])+inst .ldiv*(derw[* ,0]-derw[* ,1]) $
+(inst .phi1*(inst .delta+inst .phi1*r)+inst .phi2*inst .phi2*r)/d*(derd[* ,0]-derd[* ,1]) $
-5d-1*s*f2*inst .dis/gam*(inst .tdia*derp[* ,0]-inst .odia*derp[* ,1]))

grad=TRANPOSE(grad)
trash=WHERE(~FINITE(grad),nf)
if nf gt 0 then MESSAGE, 'Bad Jacobian.'
endif

RETURN, eff
END

```

A.8.4 Optimal estimation retrieval

```

; RACHEL nonlinear optimal estimation retrieval of extinction and backscatter
; the following need to be defined:
; a priori state (xa), covariance (sxa), structures of lidar data (rach), meteorological data (met),
; radiosonde launches (sonde), calibration function (cal), its error (scal), the time each profile
; began (start.time)
;
; switches
test=0b ; 1 - stop after first evaluation of forward model each time, 0 - don't
verbose=0b ; binary control of short or long test output
progress=0b ; 1 - plot each iteration, 0 - don't
persist=1b
ang=0b ; 1 - estimate Angstrom coefficient from AERONET, 0 - assume unity
param=0b ; 1 - calculate parameter errors, 0 - don't
ratio=0b ; 1 - retrieve lidar ratio, 0 - retrieve extinction coefficient
inter=1b

; retrieval parameters
ah=1d-2 ; scale height of vertical correlation
nex=1501 ; # bins in extinction/backscatter axis
minh=300. ; minimum retrieved height
maxh=4d3 ; maximum retrieved height
bbins=100 ; # bins used to average background
sbins=5 ; # bins used to smooth measurement variance
gamma0=1d2 ; size of first step
count.max=30 ; maximum number of iterations
eps=1d-4 ; threshold for convergence of cost
ers=10d0 ; threshold for error convergence
lr0=58. ; first guess of lidar ratio
yoff.v=6.4d1
tau.v=[16d0,16d0]
width.v=1d-2 ; GUESS

; initilias e axis
;LP.INIT,rach[0,1],el=el_,n2=n2_,range=r0
pr=WHERE(r0 gt minh AND r0 le maxh,nr)
r.y=r0[pr] ; range from instrument
r.xd=(maxh-minh)/(nex-1d0)
r.x=r.xd*DINDGEN(nex)+minh
facs=[RAYLEIGH2(rach[0,0].lambda[el_])*3/(8*!dpi),0,0,0,2d0*RAYLEIGH2(rach[0,0].lambda[el_]) $
,RAYLEIGH2(rach[0,0].lambda[el_])+RAYLEIGH2(rach[0,0].lambda[n2_]),0]

; output arrays
ny=2*nr & 1=LINDGEN(ny)
nx=2*nex+1 & m=LINDGEN(nx)
nb=2*nr+8;4*nr+6
sxa=LA.INVERT(sxa,status=sta,/double) ; inverse of a priori covariance
if sta gt 0 then $
PRINT,'Sxa inversion unstable - '+STRING(sta,format='(i0)')

; PROFILE LOOP
for ii=0,n-1 do begin;for ii=0,0 do begin ;

```

```

facs[2:3]=rach[ii].tau[[el_,n2-]]/(rach[ii].nshots*inst.width)
facs[6]=1d0/(rach[ii].nshots*inst.width)

; signal background
back=[(MOMENT(rach[ii].ch[rach[ii].nbins-bbins :*,el_] $
/(1d0-facs[2]*rach[ii].ch[rach[ii].nbins-bbins :*,el_]))], $
(MOMENT(rach[ii].ch[rach[ii].nbins-bbins :*,n2_] $
/(1d0-facs[3]*rach[ii].ch[rach[ii].nbins-bbins :*,n2_])))])

; measurement vector
y=DOUBLE(TRANPOSE([rach[ii].ch[pr,el_],rach[ii].ch[pr,n2_]])
sy=[(SMOOTH(DOUBLE(rach[ii].ch[*,el_]),sbins,/edge.trun))[pr] > 1d0, $
(SMOOTH(DOUBLE(rach[ii].ch[*,n2_]),sbins,/edge.trun))[pr] > 1d0)]
if `param then begin
  syi=DBLARR(ny,ny)
  syi[1,1]=1d0/sy
endif

; process auxilliary data
MET.INTERPOL,sonde,met,r.y,start_time[ii],rach[ii].time,num=num,ray=ray $
,v.num=num.v,v.ray=ray.v,h=inst.height,p.0=p.0,t.0=t.0
coef=1d0 & coef.v=.16d0
t0=rach[ii].lambda[el_]/rach[ii].lambda[n2_]
facs[1]=1d0+t0^coef
Df.Dcoef=(facs[1]-1d0)*ALOG(t0)

; components not in state vector
v0=rach[ii].nshots*cal[pr,1]*EXP(-facs[4]*ray)/(r.y*r.y)
v1=rach[ii].nshots*cal[pr,0]*num*EXP(-facs[5]*ray)/(r.y*r.y)

; parameter errors
sb=[tau.v,TRANPOSE(back[1],[0,1]),yoff.v,width.v,coef.v,scal[pr,0],scal[pr,1],num.v,ray.v]

; initialise loop
if ii eq 0 OR `persist $
then x = ratio ? [[REPLICATE(ALOG(1d-5),1,nex)],[REPLICATE(1r0,1,nex)],[1d0]] $
: [[REPLICATE(1d-5,1,nex)],[REPLICATE(1r0*1d-5,1,nex)],[1d0]]
count=0
gamma=gamma0
bad=0

; EVALUATE FORWARD MODEL
MODEL.PROFILE7, x, num, v0, v1, nex, nr, facs, r.x, r.y, back, f, log=ratio, ratio=ratio, k $
, Df.Dcoef, cal[pr,0], cal[pr,1], kb, inta.y=ia0
if param gt 0 then begin
  t1=MATRIX.MULTIPLY(MM(kb,sb),kb,/at) & t1[1,1]+=sy
  syi=LA.INVERT(t1,sta=sta,/dou)
  if sta gt 0 then $
    PRINT,'Sy inversion unstable - '+STRING(sta,format='(i0)')
endif

ap_cost=(x-xa)#MATRIX.MULTIPLY(sxai,x-xa,/bt)
fit_cost=(y-f)#MATRIX.MULTIPLY(syi,y-f,/bt)
cost=(ap_cost+fit_cost)[0]
if verbose then begin
  PRINT,count,ap_cost[0],fit_cost[0],cost
endif
if progress then begin
  PLOTERROR,y,1d0/SQRT(syi[1,1]),/yl,hatlen=0,background=0; ,yr=[10,1d5]
  OPLOT,f,lines=2
endif
if test then STOP

temp2=MATRIX.MULTIPLY(syi,k,/bt)
temp3=sxai + k#temp2

; save initial state
s.x=x & s.f=f & s.k=k & s.syi=syi & s.cost=cost
xs=DBLARR(nx,count.max+1) & xs[*,0]=x

; LM iteration
repeat begin
  ++count

; calculate new x
templ=((y-f)#temp2+(xa-x)#sxai)
dif = templ # LA.INVERT(gamma*sxai+temp3,sta=sta,/dou)
trash=WHERE(~FINITE(dif),ninf)
if ninf gt 0 then MESSAGE, 'Infinite dif '+STRING(ninf)
if sta gt 0 then $
  PRINT,'DIF inversion unstable - '+STRING(sta,format='(i0)')
x += dif
xs[*,count]=x

; ensure we are at a sensible state
x = (x > lower.limit) < upper.limit

MODEL.PROFILE7,x,num,v0,v1,nex,nr,facs,r.x,r.y,back, f, log=ratio, ratio=ratio, k $
,Df.Dcoef, cal[pr,0], cal[pr,1], kb, inta.y=ia0
if param then begin
  t1=MATRIX.MULTIPLY(MM(kb,sb),kb,/at) & t1[1,1]+=sy
  syi=LA.INVERT(t1,sta=sta,/dou)

```



```

    if sta gt 0 then $
        PRINT, 'Sy inversion unstable - '+STRING(sta,format='(i0)')
    endif

; calculate cost
ap_cost=(x-xa)#MATRIX.MULTIPLY(sxai,x-xa,/bt)
fit_cost=(y-f)#MATRIX.MULTIPLY(syi,y-f,/bt)
cost=(ap_cost+fit_cost)[0]
step=MATRIX.MULTIPLY(temp1,dif,/bt)
if verbose $
    then PRINT,count,ap_cost[0],fit_cost[0],cost,(s_cost-cost)/ny,step/nx
if progress then OPLOT,f,col=count/FLOAT(count_max)*253+1

; step size alterations
if cost gt s_cost then begin
    if step lt nx*eps then BREAK ;if bad ge 3 AND cost-s_cost lt ny*eps then BREAK else ++bad
    gamma /= 10d0 ; decrease step size
    x=s_x & f=s_f & k=s_k & syi=s_yi & cost=s_cost ; restore previous state
endif else begin
    bad=0
    gamma /= 2d0 ; increase step size
    temp2=MATRIX.MULTIPLY(syi,k,/bt)
    temp3=sxai + k#temp2
    sx=LA.INVERT(temp3,stat=sta,/d)
    if sta gt 0 then $
        PRINT, 'Sx inversion unstable - '+STRING(sta,format='(i0)')
    trash=WHERE(sx[m,m] lt 0.,ntrash) & if ntrash gt 0 then MESSAGE, 'Bad error matrix. Check a priori.'
    err=SQRT(sx[m,m]) ; expected standard deviation

; a point has converged if its change is less than expected error/10
; OR if the relative change in cost is below some threshold
; OR if the weighted step in x is below some threshold
if (s_cost-cost lt ny*eps) OR (TOTAL(ABS(dif) lt err/ers) eq nx) OR $
    (step lt nx*eps) then BREAK

    s_x=x & s_f=f & s_k=k & s_yi=syi & s_cost=cost ; save current state
endelse
endrep until count ge count_max
if ~verbose then PRINT,count,ap_cost[0],fit_cost[0],cost

; print out results
aver=k # syi # MATRIX.MULTIPLY(k,sx,/at)
PRINT,cost/ny,TRACE(aver)/nx

if ~param then begin
    t1=MATRIX.MULTIPLY(MM(kb,sb),kb,/at) & t1[1,1]=sy
    syi=LA.INVERT(t1,sta=sta,/dou)
    if sta gt 0 then PRINT, 'Sy inversion unstable - '+STRING(sta,format='(i0)')
    temp2=MATRIX.MULTIPLY(syi,k,/bt)
    temp3=sxai + k#temp2
    sx=LA.INVERT(temp3,stat=sta,/d)
    if sta gt 0 then PRINT, 'Sx inversion unstable - '+STRING(sta,format='(i0)')
    err=SQRT(sx[m,m]) ; expected standard deviation
endif
endif
endfor
END

```


Appendix B

Additional tables and figures

B.1 Chapter 1

The image originally presented here cannot be made freely available at the request of the copyright holder.

The image was originally published as figure 3 in Inaba and Kobayas [1972].

Figure B.1: The theoretical Raman vibrational-rotational spectrum ($\nu = 0 \rightarrow 1$ vibrational transition) at 300 K, showing the O, Q, and S-branch structures for N_2 molecule. Image reproduced from Inaba and Kobayas [1972] with kind permission from Springer Science+Business Media B.V.

B.2 Chapter 2

Table B.1: Additional details of simulated profiles for α/β retrieval

Variable	Version 1	Version 2	Version 3
Δ [mm]	2.2	2.2	
δ [mm]	1.0	2.0	
ϕ_{\parallel} [rad]	0	1.7×10^{-5}	
ϕ_{\perp} [rad]	0	0	
α_{BL} [km $^{-1}$]	0.31	0.52	0.15
α_{FT} [Mm $^{-1}$]	52	49	3
z_{BL} [km]	1.30	1.55	0.90
σ_{BL} [km]	0.1	0.08	0.2
B [sr]	41	60	45
H [km]	2.0	1.8	2.1
α_{feature} [km $^{-1}$]		0.1	60
B_{feature} [sr]		72	19
z_{feature} [km]		0.8	2.0
σ_{feature} [km]		0.15	0.08
E_L [mJ]	20	0.4	
E_B^{el} [counts]	8.7×10^{-7}	1.7×10^{-5}	
E_B^{ra} [counts]	7.1×10^{-7}	1.0×10^{-4}	

Table B.2: Dead time corrections of simulated profiles for α/β retrieval

Type	Channel	τ	P_0	P_1	P_2
Nonparalyzable	Elastic	48.7			
	Raman	58.4			
Donovan	Elastic	21.0	0.88	0.05	0.02
	Raman	27.0	0.83	0.07	0.02

Table B.3: Atmospheric parameters of simulated profiles for all retrievals

Variable	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
p_0 [hPa]	967	1130	1090	1050	1210	1220
T_0 [K]	299	283	285	270	274	281




Figure B.2: Impact of vertical correlation scale height on the logarithmic retrieval of backscatter and lidar ratio using an aerosol-derived *a priori*.




Figure B.3: Impact of vertical resolution on the logarithmic retrieval of backscatter and lidar ratio retrieval using an aerosol *a priori* and $H = 100$ m.




Figure B.4: As 2.16, but for the retrieval of extinction rather than the lidar ratio.

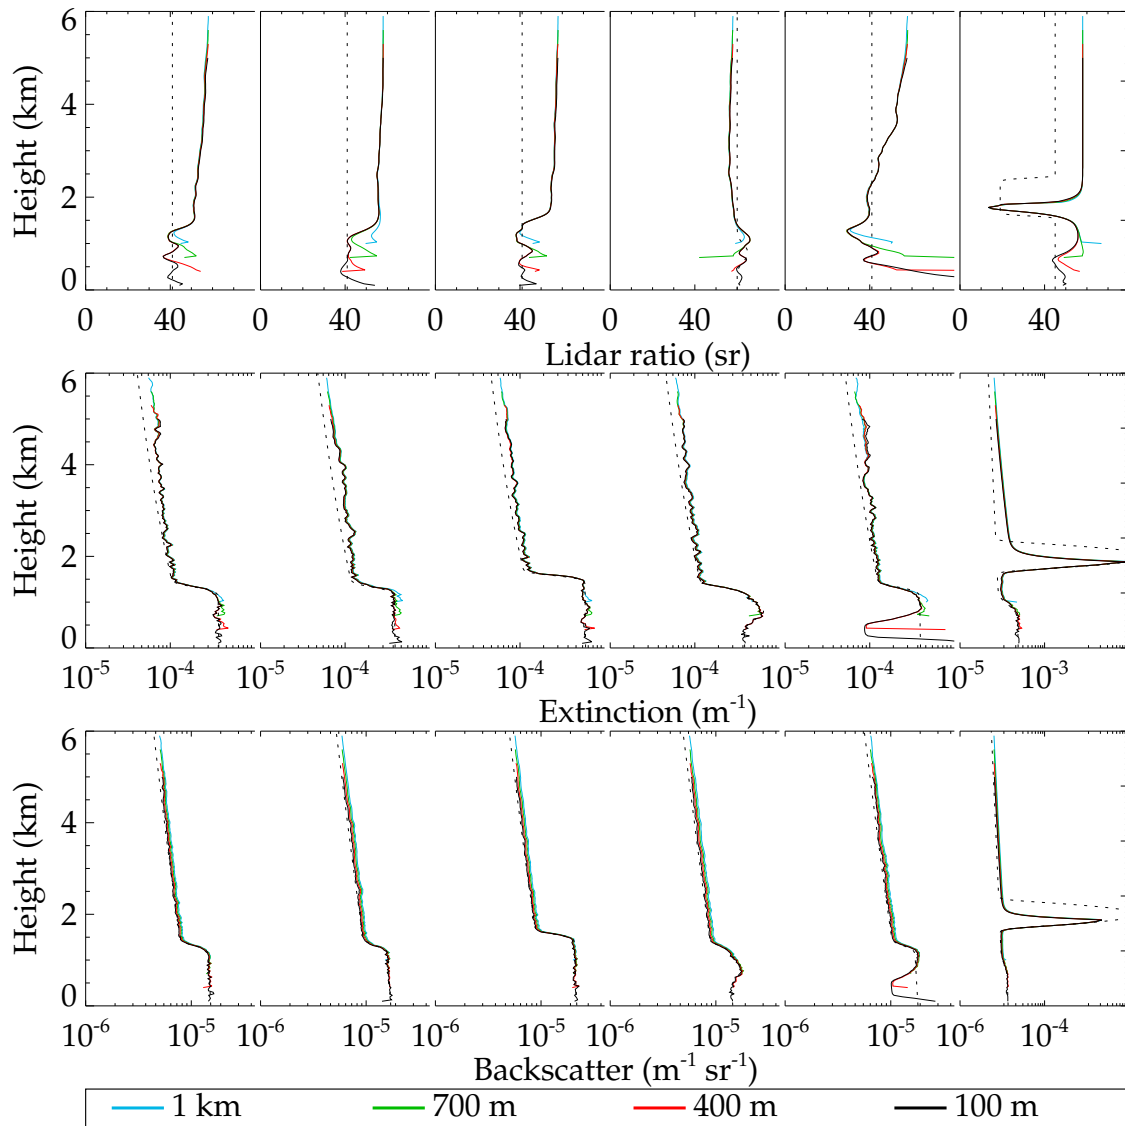


Figure B.5: Impact of minimum retrieved height on the aerosol retrieval.

B.3 Chapter 3



Figure B.6: UK Met Office radiosonde launch locations in Great Britain and Northern Ireland. (Image accessed 23 Feb 2013. Image ©2013 Google, map data ©2013 GeoBasis-DE/BKG, reproduced with permission.)

B.4 Chapter 4

Table B.4: State vectors used to simulate data for C retrieval.

Profile #	1	2	3	4	5	6
δ (mm)	12	12	12	20	8.2	7.5
ϕ_{\parallel} (mrad)	17	17	17	66	96	14
ϕ_{\perp} (μ rad)	4.1	4.1	4.1	-72	38	0.91
r_L (mm)	35	37	35	35	35	32
ϕ_L (mrad)	0.38	0.43	0.38	0.38	0.38	0.49
r_T (mm)	203	202	203	203	203	204
r_o (mm)	75	69	75	75	75	75
r_a (mm)	0.40	0.41	0.40	0.40	0.40	0.40
f (m)	2	2.01	2.00	2	2	1.99
Δ (mm)	9.3	9.3	9.3	-1.3	-3.9	-3.9
τ_L (ns)	5.2	6.1	5.2	5.2	5.2	6.4
$\eta(10^{-19} \text{ m}^2 \text{ sr}^{-1})$	1.57	1.57	1.57	1.43	1.61	2.76
z_{BL} (km)	1.3	1.3	9	1.3	1.1	1.15
H_{FT} (km)	2	2	2.3	2	1.6	1.9
α_+ (Mm^{-1})	180	180	90	180	37.0	79.0
α_- (Mm^{-1})	130	130	80	130	36.8	78.2
σ_{BL} (m)	100	100	90	100	130	80
χ_{∞}	0.503	0.503	0.176	0.503	0.0814	0.182

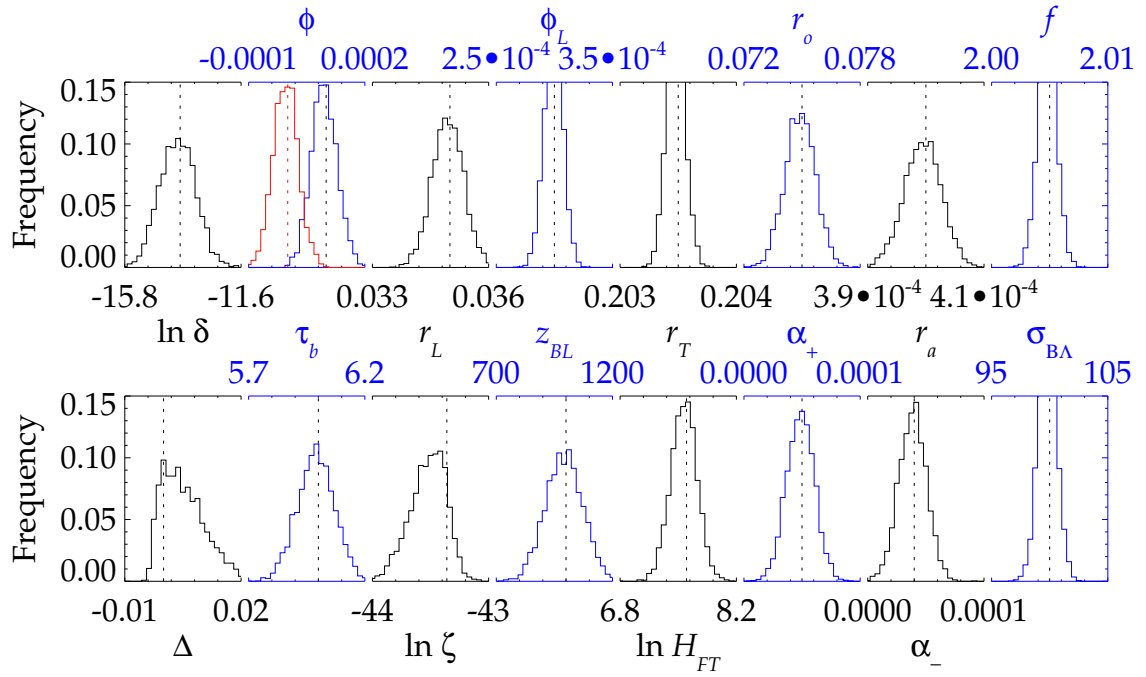


Figure B.7: As fig. 4.7, but considering $\ln \delta, \zeta, H_{FT}$.

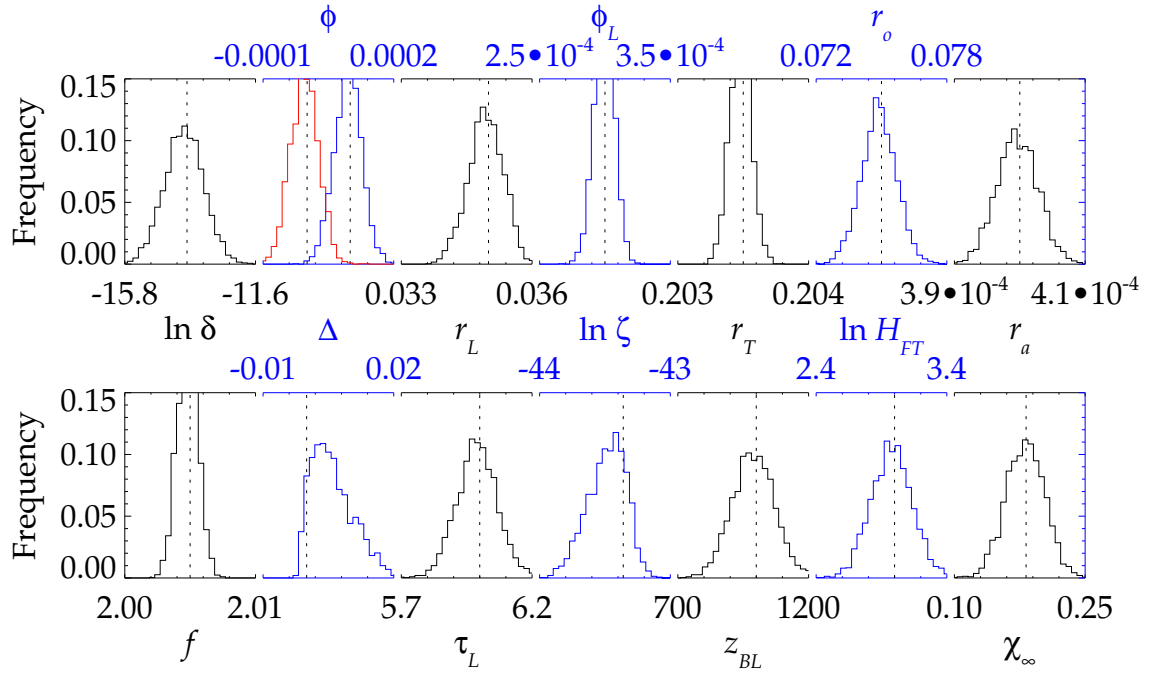


Figure B.8: As fig. 4.8, but considering $\ln \delta, \zeta, H_{FT}$.




Figure B.9: Deviation of retrieval from cases 1–7 as a function of maximum retrieved height. Line style indicates configuration.

B.5 Chapter 5




Figure B.10: Vertical positioning of a cloud, illustrated by the ABC observed by the EZ lidar (purple), CT75 ceilometer (orange), CFARR's analysis of the CUV (black), and from range-correct, raw CUV data in the analogue (green) and photon counting modes (blue).

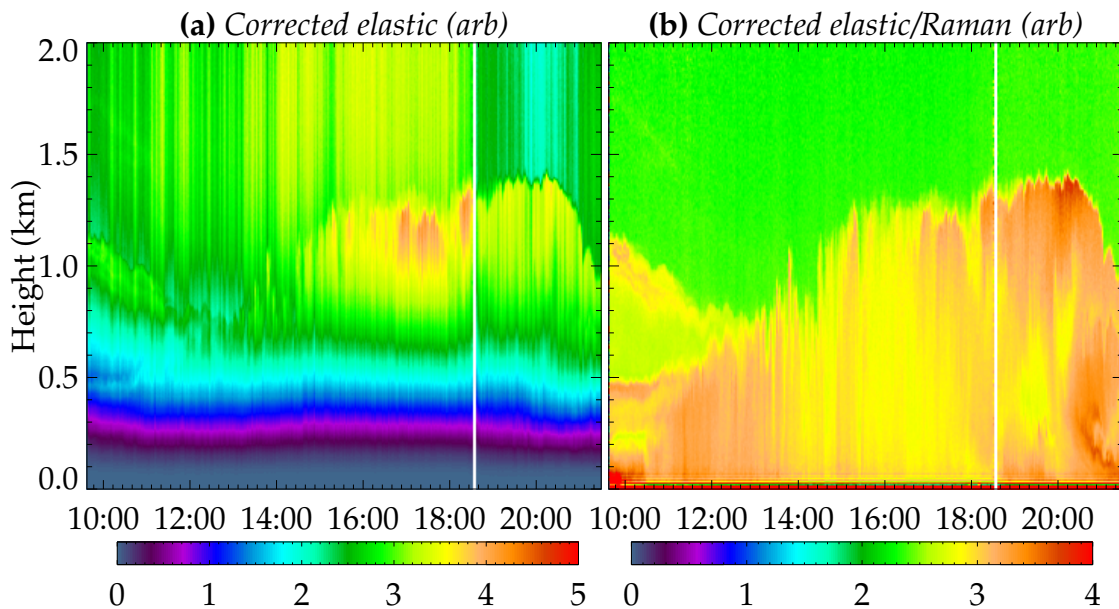


Figure B.11: Analogue observations by the CUV on 2 Mar 2010, highlighting the failures of the measured laser energy. **(a)** Elastic observations, corrected for range, background, and laser energy. **(b)** Elastic divided by Raman observations, each corrected for background.

Table B.5: Retrieved state vectors for photon counting data (see Section 5.1.3)

Profile #	1	2	3	4	5	6	7	8	9	10	11	12
δ (mm)	-34.6	-30.6	-32.4	-35.9	-32.4	28.2	29.9	-28.8	-22.1	-19.2	26.0	10.1
$\phi_{ }$ (mrad)	0.327	0.259	0.215	0.194	0.186	-0.319	-0.259	0.254	-0.0705	0.210	-0.249	0.273
ϕ_{\perp} (μ rad)	111.	-6.51	-131.	-163.	-160.	-0.00428	0.0483	0.0228	111.	0.00	-0.00221	-1.13
r_L (mm)	98.5	107.	93.0	76.2	87.5	114.	89.8	89.6	62.8	89.3	91.4	60.9
ϕ_L (mrad)	0.0491	0.0863	0.0400	0.0419	0.0400	0.133	0.194	0.183	0.0438	0.0858	0.177	0.0400
r_T (mm)	460.	460.	460.	460.	460.	460.	460.	460.	460.	460.	460.	460.
r_o (mm)	101.	123.	121.	118.	123.	116.	107.	106.	111.	98.6	106.	105.
r_a (mm)	0.942	1.08	0.820	0.839	0.793	0.699	0.756	0.727	1.07	0.753	0.735	0.993
f (m)	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00
Δ (mm)	-4.74	-2.15	-1.81	-1.89	-1.55	-8.15	-6.89	-6.80	-5.47	-6.19	-8.03	-5.18
τ_L (ns)	7.00	6.98	6.99	6.99	6.99	7.00	6.99	7.00	6.98	7.00	6.99	6.99
$\ln \eta$ ($\ln \text{ m}^2 \text{ sr}^{-1}$)	-44.2	-44.4	-43.7	-43.8	-43.7	-40.9	-41.2	-41.1	-42.1	-41.4	-41.2	-42.3
z_{BL} (km)	1.50	1.56	1.74	1.68	1.82	0.503	0.593	0.625	1.26	0.568	1.64	1.22
H_{FT} (km)	1.91	2.26	0.00100	0.00150	0.00100	2.71	3.80	3.72	1.98	3.97	3.98	3.15
α_+ (Mm^{-1})	37.1	33.8	161.	224.	173.	128.	88.2	79.5	77.5	97.2	36.9	57.1
α_- (Mm^{-1})	33.6	18.2	68.4	0.862	42.5	90.1	60.5	44.7	25.8	67.8	11.4	41.0
σ_{BL} (m)	100.	100.	100.	99.7	100.	93.6	106.	102.	101.	106.	101.	101.
χ_{∞}	0.113	0.117	0.394	0.378	0.390	0.212	0.193	0.207	0.232	0.170	0.181	0.211

Table B.6: Retrieved state vectors for analogue data (see Section 5.1.3)

Profile #	1	2	3	4	5	6	7	8	9	10	11	12
δ (mm)	-27.9	24.7	-32.2	43.0	-31.8	-28.5	-34.0	29.3	-16.2	-21.8	-12.7	-5.78
ϕ_{\parallel} (mrad)	0.231	-0.229	0.222	-0.161	0.220	0.312	0.320	-0.131	0.0843	0.250	0.0239	-0.324
ϕ_{\perp} (μ rad)	171.	0.751	41.6	159.	-27.5	8.89	4.81e-05	156.	-138.	0.00	53.2	0.772
r_L (mm)	98.3	104.	88.7	33.9	91.4	102.	92.7	71.2	71.8	89.7	88.8	72.6
ϕ_L (mrad)	0.0480	0.0692	0.0405	0.0904	0.0400	0.0993	0.158	0.282	0.0431	0.0741	0.225	0.0400
r_T (mm)	460.	460.	460.	460.	460.	460.	460.	460.	460.	460.	460.	460.
r_o (mm)	108.	114.	110.	110.	111.	105.	105.	126.	92.6	96.1	120.	115.
r_a (mm)	0.953	0.963	0.890	0.847	0.884	0.786	0.831	0.807	1.09	0.764	0.779	1.13
f (m)	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00
Δ (mm)	-3.43	-1.55	-1.58	-1.58	-1.42	-8.17	-7.67	-7.65	-5.79	-6.78	-12.3	-6.59
τ_L (ns)	6.99	7.00	6.99	6.98	6.99	6.99	6.97	7.00	6.98	6.99	6.99	6.98
$\ln \eta$ ($\ln \text{ m}^2 \text{ sr}^{-1}$)	-39.2	-39.3	-39.1	-39.0	-39.1	-37.8	-38.1	-37.9	-38.9	-38.1	-36.8	-39.1
z_{BL} (km)	0.587	0.669	1.85	1.63	1.82	0.483	0.555	0.644	1.24	0.530	1.51	1.26
H_{FT} (km)	2.01	0.504	1.94	0.0642	0.0986	1.90	2.13	3.02	1.28	3.71	2.96	2.20
α_+ (Mm $^{-1}$)	108.	98.5	116.	233.	225.	208.	135.	85.3	69.1	97.6	29.5	41.3
α_- (Mm $^{-1}$)	103.	66.4	114.	0.977	0.00100	152.	118.	48.6	26.0	72.6	6.31	35.7
σ_{BL} (m)	100.	98.6	100.	99.4	100.	96.8	134.	101.	98.3	120.	105.	102.
χ_{∞}	0.134	0.126	0.429	0.398	0.433	0.280	0.177	0.197	0.173	0.109	0.122	0.183

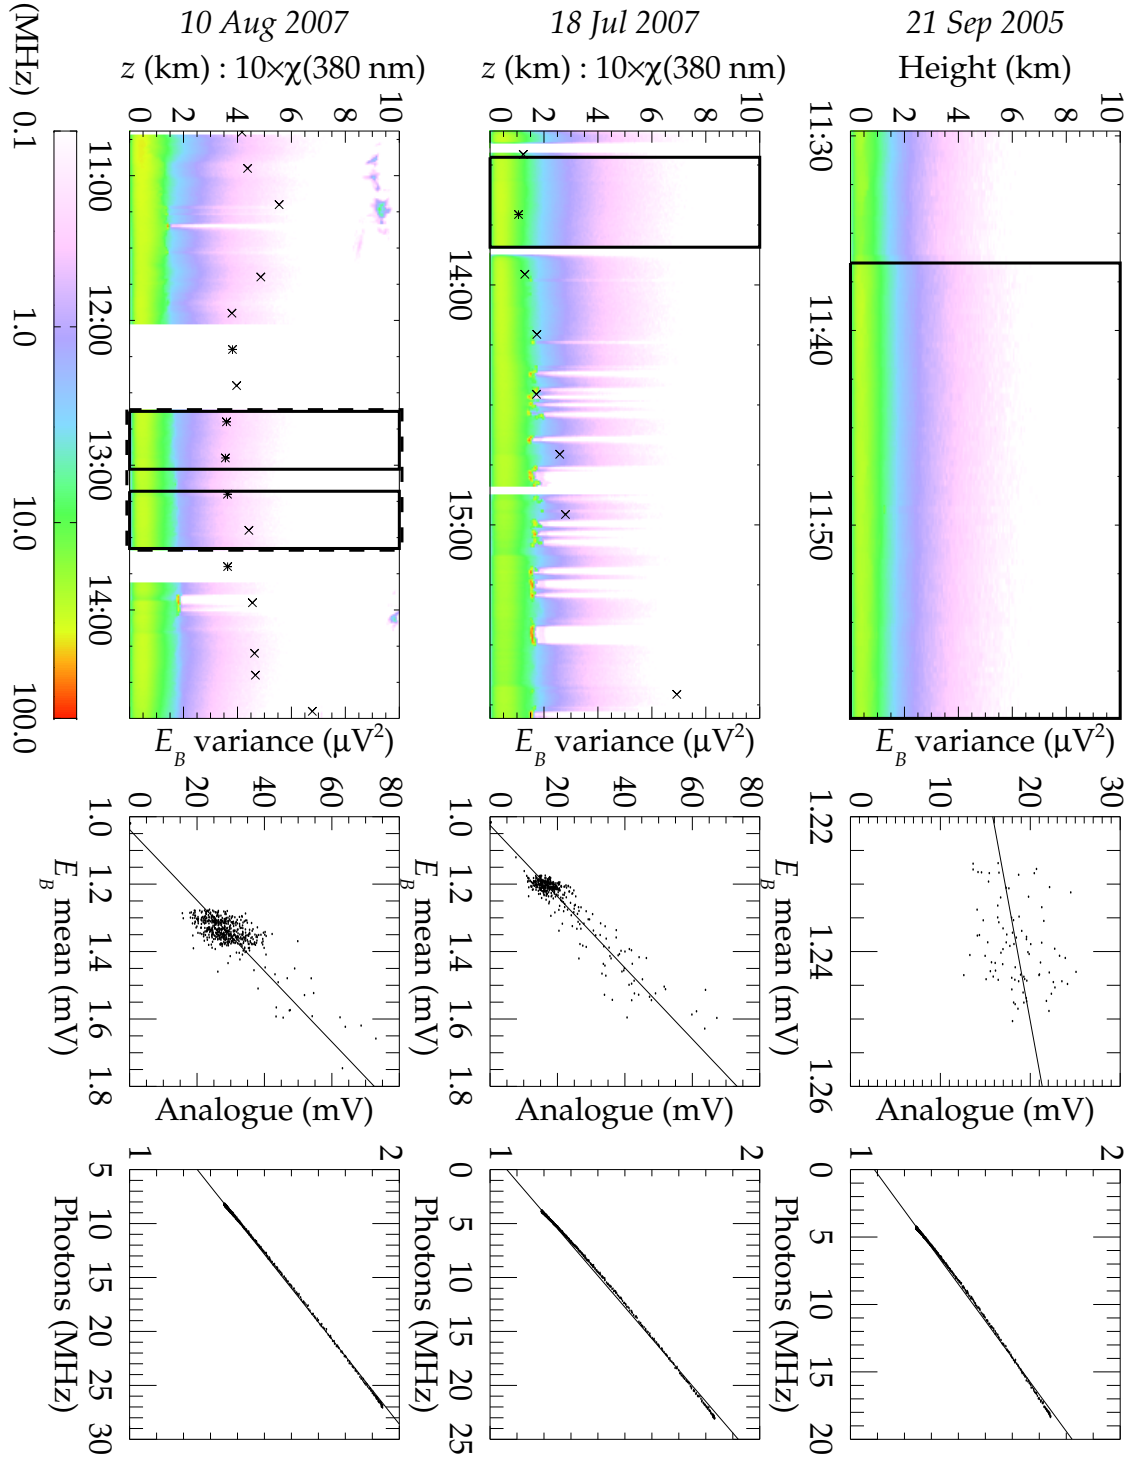


Figure B.12: The data used to produce test profiles 1–5 for the calibration function retrieval. **Left:** Elastic, PC returns on a log scale. Boxes indicate data averaged to produce each test profile. White represents missing data. $10\chi_\infty(380 \text{ nm})$ observed by AERONET is also shown with crosses for level 1.0 and stars for level 2.0 data. **Centre:** As fig. 5.2. **Right:** Analogue vs. dead time-corrected PC returns of the averaged profile.

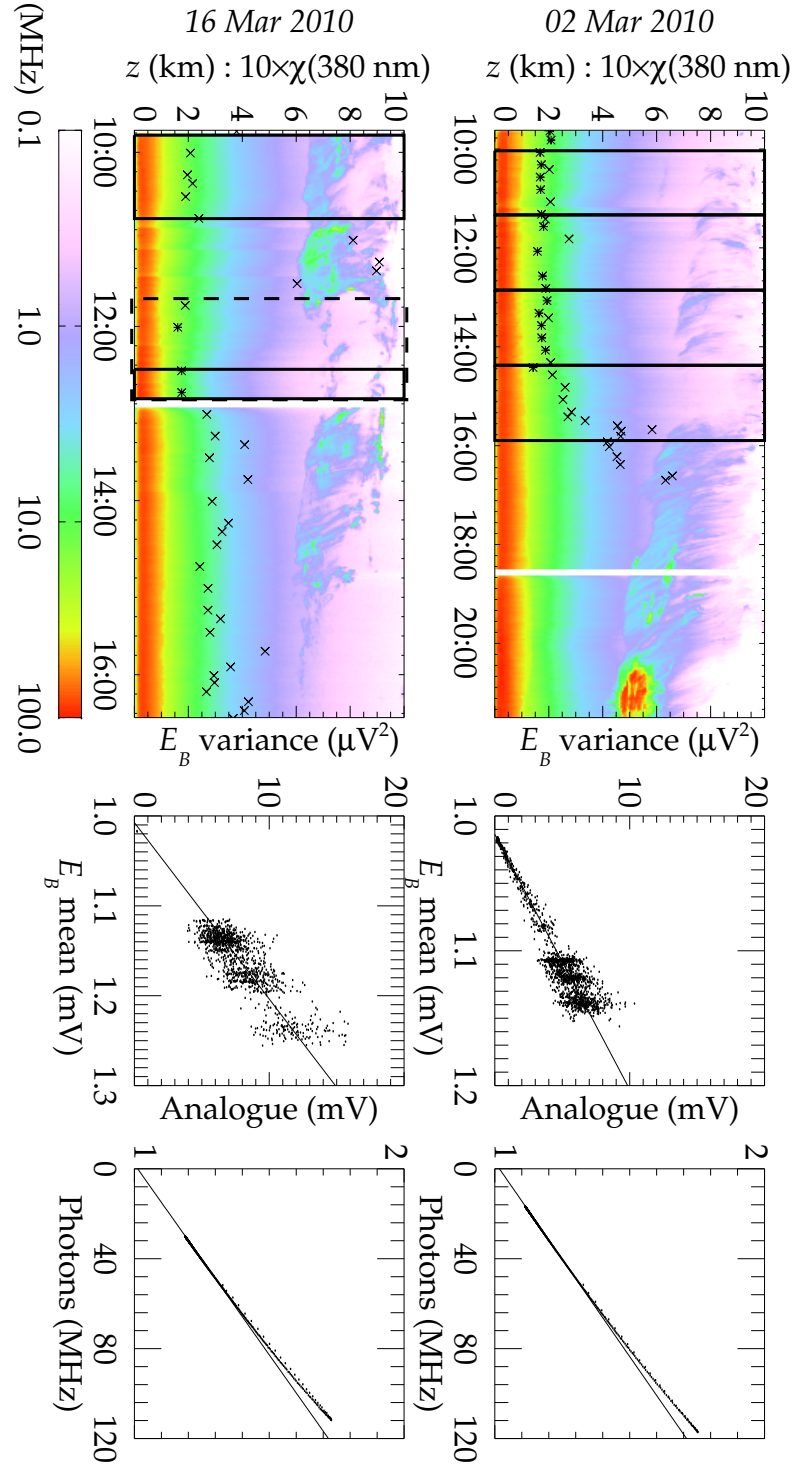


Figure B.13: As fig. B.12, but for test profiles 6–12. Note the failure of the nonlinear correction at high count rates.

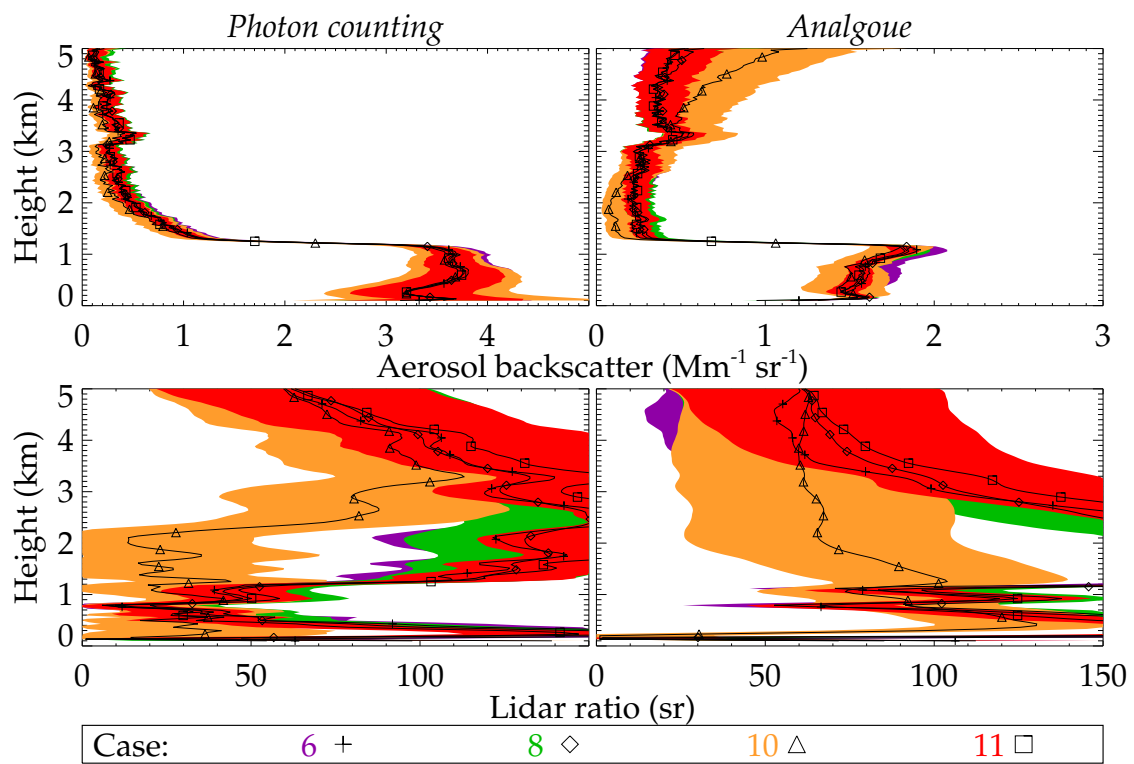


Figure B.14: As fig. 5.16, but showing the retrieved errors for the solutions in colour and the solution with lines. For clarity, only cases 6, 8, 10, and 11 are presented.

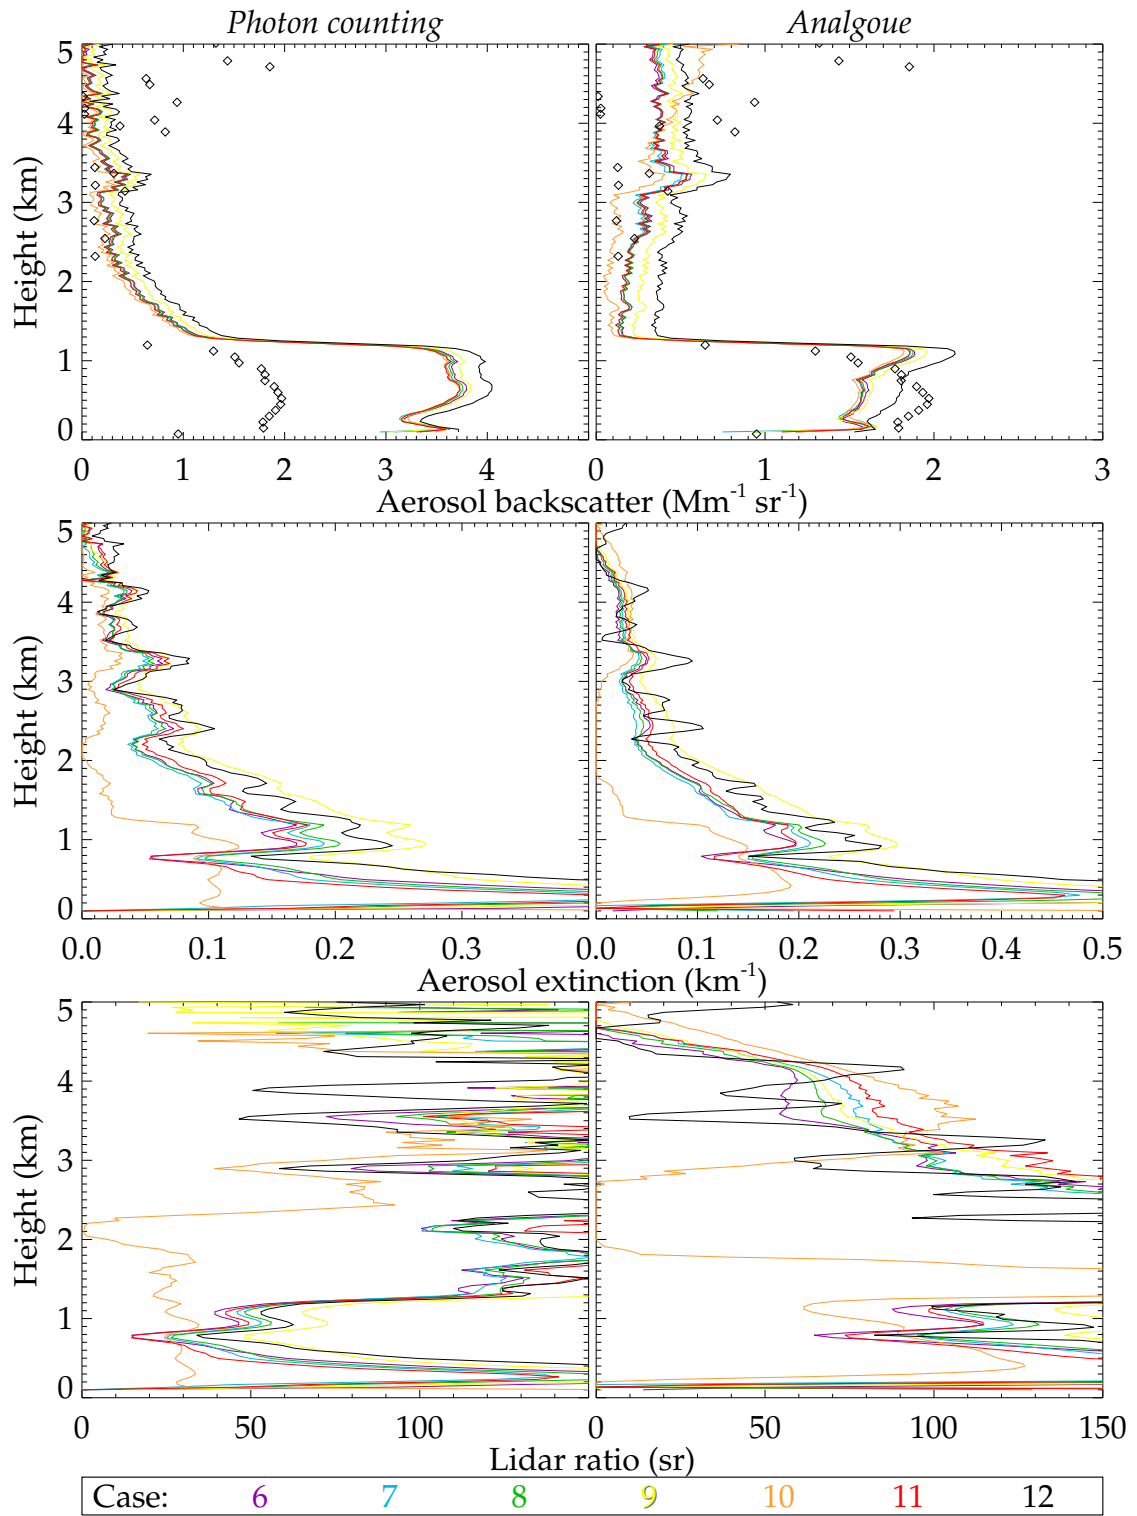


Figure B.15: As fig. 5.16, but for the linear retrieval configuration.

Bibliography

- Adam, M., Kovalev, V. A., Wold, C., Newton, J., Pahlow, M., Hao, W. M., and Parlange, M. B. (2007). Application of the Kano-Hamilton multiangle inversion method in clear atmospheres. *Journal of Atmospheric and Oceanic Technology*, 24(12):2014–2028, doi:10.1175/2007jtecha946.1.
- Agnew, J. and Wrench, C. (2006 – 2010). Chilbolton UV Raman lidar raw data. STFC Chilbolton Observatory, Rutherford Appleton Laboratory.
- Agnew, J. L. (2002). Lidar measurement of tropospheric radio refractivity. Technical report, Rutherford-Appleton Laboratory. Available at <http://www.ofcom.org.uk/static/archive/ra/topics/research/rcru/project41/finalreport/contents.htm>.
- Agnew, J. L. (2003). Lidar and radar tropospheric profiling at Chilbolton Observatory. During *Sixth International Symposium on Tropospheric Profiling: Needs and Technologies*, pages 151–153.
- Althausen, D., Müller, D., Ansmann, A., Wandinger, U., Hube, H., Clauder, E., and Zörner, S. (2000). Scanning 6-wavelength 11-channel aerosol lidar. *Journal of Atmospheric and Oceanic Technology*, 17(11):1469–1482, doi:10.1175/1520-0426(2000)017<1469:swcal>2.0.co;2.
- Ancellet, G. M., Kavaya, M. J., Menzies, R. T., and Brothers, A. M. (1986). Lidar telescope overlap function and effects of misalignment for unstable resonator transmitter and coherent receiver. *Applied Optics*, 25(17):2886–2890, doi:10.1364/AO.25.002886.
- Ansmann, A., Riebesell, M., and Weitkamp, C. (1990). Measurement of atmospheric aerosol extinction profiles with a Raman lidar. *Optics Letters*, 15(13):746–748, doi:10.1364/OL.15.000746.

- Ansmann, A., Tesche, M., Gross, S., Freudenthaler, V., Seifert, P., Hiebsch, A., Schmidt, J., Wandinger, U., Mattis, I., Müller, D., and Wiegner, M. (2010). The 16 April 2010 major volcanic ash plume over central Europe: EARLINET lidar and AERONET photometer observations at Leipzig and Munich, Germany. *Geophysical Research Letters*, 37:L13810, doi:10.1029/2010gl043809.
- Ansmann, A., Tesche, M., Seifert, P., Gross, S., Freudenthaler, V., Apituley, A., Wilson, K. M., Serikov, I., Linné, H., Heinold, B., Hiebsch, A., Schnell, F., Schmidt, J., Mattis, I., Wandinger, U., and Wiegner, M. (2011). Ash and fine-mode particle mass profiles from EARLINET-AERONET observations over central Europe after the eruptions of the Eyjafjallajökull volcano in 2010. *Journal of Geophysical Research — Atmospheres*, 116:D00U02, doi:10.1029/2010jd015567.
- Ansmann, A., Wandinger, U., Riebesell, M., Weitkamp, C., and Michaelis, W. (1992). Independent measurement of extinction and backscatter profiles in cirrus clouds by using a combined Raman elastic-backscatter lidar. *Applied Optics*, 31(33):7113–7131, doi:10.1364/AO.31.007113.
- Argall, P. and Sica, R. (2003). *Lidar: Atmospheric sounding introduction*, pages 1169–1176. Academic Press, San Diego, CA.
- Barlow, J. F., Dunbar, T. M., Nemitz, E. G., Wood, C. R., Gallagher, M. W., Davies, F., O'Connor, E., and Harrison, R. M. (2011). Boundary layer dynamics over London, UK, as observed using Doppler lidar during REPARTEE-II. *Atmospheric Chemistry and Physics*, 11(5):2111–2125, doi:10.5194/acp-11-2111-2011.
- Berezhnyy, I. (2009). A combined diffraction and geometrical optics approach for lidar overlap function computation. *Optics and Lasers in Engineering*, 47(7-8):855–859, doi:10.1016/j.optlaseng.2009.01.011.
- Berkoff, T. A., Welton, E. J., Campbell, J. R., Scott, V. S., and Spinhirine, J. D. (2003). Investigation of overlap correction techniques for the micro-pulse lidar network (MPLNET). During *IGARSS 2003: IEEE International Geoscience and Remote Sensing Symposium*, pages 4395–4397.
- Biavati, G., Di Donfrancesco, G., Cairo, F., and Feist, D. G. (2011). Correction scheme for close-range lidar returns. *Applied Optics*, 50(30):5872–5882, doi:10.1364/AO.50.005872.

- Bissonnette, L. R. (1986). Sensitivity analysis of lidar inversion algorithms. *Applied Optics*, 25(13):2122–2125, doi:10.1364/AO.25.002122.
- Böckmann, C. (2001). Hybrid regularization method for the ill-posed inversion of multiwavelength lidar data in the retrieval of aerosol size distributions. *Applied Optics*, 40(9):1329–1342, doi:10.1364/AO.40.001329.
- Böckmann, C., Wandinger, U., Ansmann, A., Bösenberg, J., Amiridis, V., Boselli, A., Delaval, A., De Tomasi, F., Frioud, M., Grigorov, I. V., Hågård, A., Horvat, M., Iarlori, M., Komguem, L., Kreipl, S., Larchêvque, G., Matthias, V., Papayannis, A., Pappalardo, G., Rocadenbosch, F., Rodrigues, J. A., Schneider, J., Shcherbakov, V., and Wiegner, M. (2004). Aerosol lidar intercomparison in the framework of the EARLINET project. 2. Aerosol backscatter algorithms. *Applied Optics*, 43(4):977–989, doi:10.1364/AO.43.000977.
- Bohren, C. F. and Huffman, D. R. (2004). *Absorption and Scattering of Light by Small Particles*. Wiley-VCH, Mörlenbach, Germany.
- Bradley, S., Perrott, Y., and Behrens, P. (2010). Myres Hill LIDAR/SODAR/Mast intercomparison revisited. During *15th International Symposium for the Advancement of Remote Sensing of the Boundary Layer*. http://www.isars2010.uvvsq.fr/images/stories/Presentations/O_MEA_08_Bradley.ppt.
- Brooks, I. M. (2003). Finding boundary layer top: Application of a wavelet covariance transform to lidar backscatter profiles. *Journal of Atmospheric and Oceanic Technology*, 20(8):1092–1105, doi:10.1175/1520-0426(2003)020<1092:FBLTAO>2.0.CO;2.
- Bucholtz, A. (1995). Rayleigh-scattering calculations for the terrestrial atmosphere. *Applied Optics*, 34(15):2765–2773, doi:10.1364/AO.34.002765.
- Cao, X., Roy, G., Roy, N., and Bernier, R. (2009). Comparison of the relationships between lidar integrated backscattered light and accumulated depolarization ratios for linear and circular polarization for water droplets, fog oil, and dust. *Applied Optics*, 48(21):4130–4141, doi:10.1364/AO.48.004130.
- Chib, S. and Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. *American Statistician*, 49(4):327–335, doi:10.2307/2684568.

- Colbeck, I. (1998). *Physical and Chemical Properties of Aerosols*. Blackie Academic and Professional, London, first edition.
- Collis, R. T. H. (1966). Lidar — a new atmospheric probe. *Quarterly Journal of the Royal Meteorological Society*, 92(392):220–230, doi:10.1002/qj.49709239205.
- Dacre, H. F., Grant, A. L. M., Hogan, R. J., Belcher, S. E., Thomson, D. J., Devenish, B. J., Marengo, F., Hort, M. C., Haywood, J. M., Ansmann, A., Mattis, I., and Clarisse, L. (2011). Evaluating the structure and magnitude of the ash plume during the initial phase of the 2010 Eyjafjallajökull eruption using lidar observations and NAME simulations. *Journal of Geophysical Research — Atmospheres*, 116:D00U03, doi:10.1029/2011jd015608.
- Delanoe, J. and Hogan, R. J. (2008). A variational scheme for retrieving ice cloud properties from combined radar, lidar, and infrared radiometer. *Journal of Geophysical Research — Atmospheres*, 113:D07204, doi:10.1029/2007jd009000.
- Dho, S. W., Park, Y. J., and Kong, H. J. (1997). Experimental determination of a geometric form factor in a lidar equation for an inhomogeneous atmosphere. *Applied Optics*, 36(24):6009–6010, doi:10.1364/ao.36.006009.
- Di Girolamo, P., Ambrico, P. F., Amodeo, A., Boselli, A., Pappalardo, G., and Spinelli, N. (1999). Aerosol observations by lidar in the nocturnal boundary layer. *Applied Optics*, 38(21):4585–4595, doi:10.1364/AO.38.004585.
- Donovan, D. P., Whiteway, J. A., and Carswell, A. I. (1993). Correction for nonlinear photon-counting effects in lidar systems. *Applied Optics*, 32(33):6742–6753, doi:10.1364/AO.32.006742.
- Dubovik, O., Sinyuk, A., Lapyonok, T., Holben, B. N., Mishchenko, M., Yang, P., Eck, T. F., Volten, H., Munoz, O., Veihelmann, B., van der Zande, W. J., Leon, J. F., Sorokin, M., and Slutsker, I. (2006). Application of spheroid models to account for aerosol particle nonsphericity in remote sensing of desert dust. *Journal of Geophysical Research — Atmospheres*, 111:D11208, doi:10.1029/2005jd006619.
- Electronic Tube Centre (ETC) (1998). Metal package photomultiplier tubes: R7400U series and subminiature photosensor modules. Technical Report TPMH9003E01, Hamamatsu Photonics K.K.

- Eloranta, E. W. (1998). Practical model for the calculation of multiply scattered lidar returns. *Applied Optics*, 37(12):2464–2472, doi:10.1364/AO.37.002464.
- Evans, R. D. (1955). *The Atomic Nucleus*, chapter 28. McGraw-Hill, New Delhi. A copy can be found at <http://archive.org/details/atomicnucleus032805mbp>.
- Ewell, D. M., Flocchini, R. G., Myrup, L. O., and Cahill, T. A. (1989). Aerosol transport in the southern Sierra-Nevada. *Journal of Applied Meteorology*, 28(2):112–125, doi:10.1175/1520-0450(1989)028<0112:atitss>2.0.co;2.
- Fernald, F., Herman, B., and Reagan, J. (1972). Determination of aerosol height distributions by lidar. *Journal of Applied Meteorology*, 11(3):482–489, doi:10.1175/1520-0450(1972)011<0482:DOAHDB>2.0.CO;2.
- Fernald, F. G. (1984). Analysis of atmospheric lidar observations — some comments. *Applied Optics*, 23(5):652–653, doi:10.1364/AO.23.000652.
- Ferrare, R. A., Melfi, S. H., Whiteman, D. N., Evans, K. D., and Leifer, R. (1998). Raman lidar measurements of aerosol extinction and backscattering. 1. Methods and comparisons. *Journal of Geophysical Research — Atmospheres*, 103(D16):19663–19672, doi:10.1029/98jd01646.
- Ferrare, R. A., Turner, D. D., Brasseur, L. H., Feltz, W. F., Dubovik, O., and Tooman, T. P. (2001). Raman lidar measurements of the aerosol extinction-to-backscatter ratio over the Southern Great Plains. *Journal of Geophysical Research — Atmospheres*, 106(D17):20333–20347, doi:10.1029/2000JD000144.
- Ferrero, L., Mocnik, G., Ferrini, B. S., Perrone, M. G., Sangiorgi, G., and Bolzacchini, E. (2011). Vertical profiles of aerosol absorption coefficient from micro-aethalometer data and Mie calculation over Milan. *Science of the Total Environment*, 409(14):2824–2837, doi:10.1016/j.scitotenv.2011.04.022.
- Ferrero, L., Perrone, M. G., Petraccone, S., Sangiorgi, G., Ferrini, B. S., Lo Porto, C., Lazzati, Z., Cocchi, D., Bruno, F., Greco, F., Riccio, A., and Bolzacchini, E. (2010). Vertically-resolved particle size distribution within and above the mixing layer over the Milan metropolitan area. *Atmospheric Chemistry and Physics*, 10(8), doi:10.5194/acp-10-3915-2010.

- Fried, D. L. (1965). Noise in photoemission current. *Applied Optics*, 4(1):79–80, doi:10.1364/AO.4.000079.
- Friedland, S. S., Katzenstein, J., and Zatzick, M. R. (1956). Pulsed search-lighting the atmosphere. *Journal of Geophysical Research*, 61(3):415–434, doi:10.1029/JZ061i003p00415.
- Fugii, T. and Fukuchi, T. (2005). *Laser Remote Sensing*. Taylor and Francis, Boca Raton, FL.
- Ghate, V. P., Albrecht, B. A., and Kollias, P. (2010). Vertical velocity structure of nonprecipitating continental boundary layer stratocumulus clouds. *Journal of Geophysical Research — Atmospheres*, 115:D13204, doi:10.1029/2009jd013091.
- Goldsmith, J. E. M., Blair, F. H., Bisson, S. E., and Turner, D. D. (1998). Turn-key Raman lidar for profiling atmospheric water vapor, clouds, and aerosols. *Applied Optics*, 37(21):4979–4990, doi:10.1364/AO.37.004979.
- Grainger, R. G., Lucas, J., Thomas, G. E., and Ewen, G. B. L. (2004). Calculation of Mie derivatives. *Applied Optics*, 43(28):5386–5393, doi:10.1364/ao.43.005386. Codes may be found at <http://www-atm.physics.ox.ac.uk/code/mie/> (last accessed 21 Jul 2013).
- Granados-Munoz, M. J., Navas-Guzman, F., Bravo-Aranda, J. A., Guerrero-Rascado, J. L., Lyamani, H., Fernandez-Galvez, J., and Alados-Arboledas, L. (2012). Automatic determination of the planetary boundary layer height using lidar: One-year analysis over southeastern Spain. *Journal of Geophysical Research — Atmospheres*, 117:D18208, doi:10.1029/2012jd017524.
- Grant, W. B. (1991). Differential absorption and Raman lidar for water-vapor profile measurements — A review. *Optical Engineering*, 30(1):40–48, doi:10.1117/12.55772.
- Greenberg, J. R., Guenther, A. B., and Turnipseed, A. (2009). Tethered balloon-based soundings of ozone, aerosols, and solar radiation near Mexico City during MIRAGE-MEX. *Atmospheric Environment*, 43(16):2672–2677, doi:10.1016/j.atmosenv.2009.02.019.
- Guerrero-Rascado, J. L., Costa, M. J., Bortoli, D., Silva, A. M., Lyamani, H., and Alados-Arboledas, L. (2010). Infrared lidar overlap function: an experimental determination. *Optics Express*, 18(19):20350–20359, doi:10.1364/OE.18.020350.

- Guerrero-Rascado, J. L., Ruiz, B., Chourdakis, G., Georgoussis, G., and Alados-Arboledas, L. (2008). One year of water vapour Raman lidar measurements at the Andalusian Centre for Environmental Studies (CEAMA). *International Journal of Remote Sensing*, 30:5437–5453, doi:10.1080/01431160802036433.
- Guldner, J. and Spankuch, D. (2001). Remote sensing of the thermodynamic state of the atmospheric boundary layer by ground-based microwave radiometry. *Journal of Atmospheric and Oceanic Technology*, 18(6):925–933, doi:10.1175/1520-0426(2001)018<0925:rsotts>2.0.co;2.
- Haefele, A., Barras, E. M., Maier, O., Ruffieus, D., and Calpini, B. (2012). Composite temperature profiles from Raman lidar and microwave radiometer. During *Ninth International Symposium on Tropospheric Profiling*. http://cetemps.aquila.infn.it/istp/proceedings/Session_E_Validation_Instruments_Synergies_and_Field_Experime/SE_10_Haefele.pdf.
- Haefelin, M., Angelini, F., Morille, Y., Martucci, G., Frey, S., Gobbi, G. P., Lolli, S., O'Dowd, C. D., Sauvage, L., Xueref-Remy, I., Wastine, B., and Feist, D. G. (2012). Evaluation of mixing-height retrievals from automatic profiling lidars and ceilometers in view of future integrated networks in Europe. *Boundary-Layer Meteorology*, 143(1):49–75, doi:10.1007/s10546-011-9643-z.
- Hair, J. W., Hostetler, C. A., Cook, A. L., Harper, D. B., Ferrare, R. A., Mack, T. L., Welch, W., Izquierdo, L. R., and Hovis, F. E. (2008). Airborne high spectral resolution lidar for profiling aerosol optical properties. *Applied Optics*, 47(36):6734–6752, doi:10.1364/ao.47.006734.
- Hakamata, T. (2006). Photomultiplier tubes: Basics and applications. Technical report, Hamamatsu Photonics K.K. third edition.
- Halldórsson, T. and Langerholc, J. (1978). Geometrical form factors for the lidar function. *Applied Optics*, 17(2):240–244, doi:10.1364/AO.17.000240.
- Hansen, P. C. (1992). Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34(4):561–580, doi:10.1137/1034115.
- Haywood, J. M. and Shine, K. P. (1995). The effect of anthropogenic sulfate and soot aerosol on the clear-sky planetary radiation budget. *Geophysical Research Letters*, 22(5):603–606, doi:10.1029/95gl00075.

- He, Q. S., Li, C. C., Mao, J. T., Lau, A. K. H., and Chu, D. A. (2008). Analysis of aerosol vertical distribution and variability in Hong Kong. *Journal of Geophysical Research — Atmospheres*, 113:D14211, doi:10.1029/2008jd009778.
- Hervo, M., Quennehen, B., Kristiansen, N. I., Boulon, J., Stohl, A., Fréville, P., Pichon, J. M., Picard, D., Labazuy, P., Gouhier, M., Roger, J. C., Colomb, A., Schwarzenboeck, A., and Sellegri, K. (2012). Physical and optical properties of 2010 Eyjafjallajökull volcanic eruption aerosol: ground-based, lidar and airborne measurements in France. *Atmospheric Chemistry and Physics*, 12(4):1721–1736, doi:10.5194/acp-12-1721-2012.
- Hess, M., Koepke, P., and Schult, I. (1998). Optical properties of aerosols and clouds: The software package OPAC. *Bulletin of the American Meteorological Society*, 79(5):831–844, doi:10.1175/1520-0477(1998)079<0831:opoaac>2.0.co;2.
- Hey, J. V., Coupland, J., Foo, M. H., Richards, J., and Sandford, A. (2011). Determination of overlap in lidar systems. *Applied Optics*, 50(30):5791–7, doi:10.1364/ao.50.005791.
- Hinds, W. (1982). *Aerosol Technology*. Wiley, New York.
- Hogan, R. J., Grant, A. L. M., Illingworth, A. J., Pearson, G. N., and O'Connor, E. J. (2009). Vertical velocity variance and skewness in clear and cloud-topped boundary layers as revealed by Doppler lidar. *Quarterly Journal of the Royal Meteorological Society*, 135(640):635–643, doi:10.1002/qj.413.
- Holben, B. N., Eck, T. F., Slutsker, I., Tanre, D., Buis, J. P., Setzer, A., Vermote, E., Reagan, J. A., Kaufman, Y. J., Nakajima, T., Lavenu, F., Jankowiak, I., and Smirnov, A. (1998). AERONET - a federated instrument network and data archive for aerosol characterization. *Remote Sensing of Environment*, 66(1):1–16, doi:10.1016/s0034-4257(98)00031-5.
- Huang, Z. W., Huang, J. P., Bi, J. R., Wang, G. Y., Wang, W. C., Fu, Q. A., Li, Z. Q., Tsay, S. C., and Shi, J. S. (2010). Dust aerosol vertical structure measurements using three MPL lidars during 2008 China-US joint dust field experiment. *Journal of Geophysical Research — Atmospheres*, 115:D00K15, doi:10.1029/2009jd013273.
- Hulburt, E. O. (1937). Observations of a searchlight beam to an altitude of 28 kilometers. *Journal of the Optical Society of America*, 27(11):377–382, doi:10.1364/JOSA.27.000377.

- Inaba, H. and Kobayas, T. (1972). Laser-Raman radar — Laser-Raman scattering methods for remote detection and analysis of atmospheric pollution. *Opto-Electronics*, 4(2):101–123, doi:10.1007/bf01421175.
- Intergovernmental Panel on Climate Change (IPCC), Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K. B., Tignor, M., and Miller, H. L. (2007). *Climate Change 2007: The Physical Basis of Climate Change. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, New York.
- Johnson, B., Turnbull, K., Brown, P., Burgess, R., Dorsey, J., Baran, A. J., Webster, H., Haywood, J., Cotton, R., Ulanowski, Z., Hesse, E., Woolley, A., and Rosenberg, P. (2012). In situ observations of volcanic ash clouds from the FAAM aircraft during the eruption of Eyjafjallajökull in 2010. *Journal of Geophysical Research — Atmospheres*, 117:D00U24, doi:10.1029/2011jd016760.
- King, L. (1923). On the complex anisotropic molecule in relation to the dispersion and scattering of light. *Proceedings of the Royal Society A*, 104:333–357, doi:10.1098/rspa.1923.0113.
- Kitchen, M. (1989). Representativeness errors for radiosonde observations. *Quarterly Journal of the Royal Meteorological Society*, 115(487):673–700, doi:10.1256/smsqj.48712.
- Klett, J. D. (1981). Stable analytical inversion solution for processing lidar returns. *Applied Optics*, 20(2):211–220, doi:10.1364/AO.20.000211.
- Klett, J. D. (1985). Lidar inversion with variable backscatter extinction ratios. *Applied Optics*, 24(11):1638–1643, doi:10.1364/AO.24.001638.
- Kuang, S., Burris, J. F., Newchurch, M. J., Johnson, S., and Long, S. (2011). Differential absorption lidar to measure subhourly variation of tropospheric ozone profiles. *IEEE Transactions on Geoscience and Remote Sensing*, 49(1):557–571, doi:10.1109/tgrs.2010.2054834.
- Kuze, H., Kinjo, H., Sakurada, Y., and Takeuchi, N. (1998). Field-of-view dependence of lidar signals by use of Newtonian and Cassegrainian telescopes. *Applied Optics*, 37(15):3128–3132, doi:10.1364/AO.37.003128.

- Li, W., Stamnes, K., Spurr, R., and Stamnes, J. (2008). Simultaneous retrieval of aerosol and ocean properties by optimal estimation: SeaWiFS case studies for the Santa Barbara Channel. *International Journal of Remote Sensing*, 29(19):5689–5698, doi:10.1080/01431160802007632.
- Lohmann, U. and Feichter, J. (2005). Global indirect aerosol effects: a review. *Atmospheric Chemistry and Physics*, 5(3):715–737, doi:10.5194/acp-5-715-2005.
- Mao, F. Y., Gong, W., and Li, J. (2012). Geometrical form factor calculation using Monte Carlo integration for lidar. *Optics and Laser Technology*, 44(4):907–912, doi:10.1016/j.optlastec.2011.10.024.
- Marchant, C. C., Moon, T. K., and Gunther, J. H. (2010). An iterative least square approach to elastic-lidar retrievals for well-characterized aerosols. *IEEE Transactions on Geoscience and Remote Sensing*, 48(5):2430–2444, doi:10.1109/tgrs.2009.2038903.
- Marchant, C. C., Wojcik, M. D., and Bradford, W. J. (2012). Estimation of aerosol effective radius by multiwavelength elastic lidar. *IEEE Transactions on Geoscience and Remote Sensing*, 50(2):645–660, doi:10.1109/tgrs.2011.2160725.
- Marenco, F. and Hogan, R. J. (2011). Determining the contribution of volcanic ash and boundary layer aerosol in backscatter lidar returns: A three-component atmosphere approach. *Journal of Geophysical Research — Atmospheres*, 116:D00U06, doi:10.1029/2010jd015415.
- Marks, C. J. and Rodgers, C. D. (1993). A retrieval method for atmospheric composition from limb emission measurements. *Journal of Geophysical Research — Atmospheres*, 98(D8):14939–14953, doi:10.1029/93jd01195.
- Mattis, I., Flentje, H., and Thomas, W. (2012). The DWD ceilometer network for aerosol profiling. During *Ninth International Symposium on Tropospheric Profiling*. http://cetemps.aquila.infn.it/istp/proceedings/Session_G_Networks_Satellites_and_Aircrafts/SG_03_Mattis.pdf.
- Maurya, R., Dubey, P. K., Arya, B. C., and Jain, S. L. (2010). Aerosol vertical profile measurements using micro pulse lidar at New Delhi, India. During *15th International Symposium for the Advancement of Remote Sensing of the Boundary Layer*. http://www.isars2010.uvsq.fr/images/stories/Presentations/O_MEA_04_Maurya.pdf.

- Measures, R. M. (1992). *Lidar Remote Sensing: Fundamentals and Applications*. Krieger Publishing Company, Malabar, Florida, second edition.
- Melfi, S. H. (1972). Remote measurements of the atmosphere using Raman scattering. *Applied Optics*, 11(7):1605–1610, doi:10.1364/AO.11.001605.
- Mie, G. (1908). Beiträge zur optik trüber medien, speziell kolloidaler metallösungen. *Annalen der Physik*, 25(3):377–445. A translation can be found at <http://diogenes.iwt.uni-bremen.de/vt/laser/papers/RAE-LT1873-1976-Mie-1908-translation.pdf>.
- Morille, Y., Haeffelin, M., Drobinski, P., and Pelon, J. (2007). STRAT: An automated algorithm to retrieve the vertical structure of the atmosphere from single-channel lidar data. *Journal of Atmospheric and Oceanic Technology*, 24(5):761–775, doi:10.1175/jtech2008.1.
- Müller, D., Ansmann, A., Mattis, I., Tesche, M., Wandinger, U., Althausen, D., and Pisani, G. (2007). Aerosol-type-dependent lidar ratios observed with Raman lidar. *Journal of Geophysical Research — Atmospheres*, 112:D16202, doi:10.1029/2006jd008292.
- Müller, D., Wandinger, U., and Ansmann, A. (1999). Microphysical particle parameters from extinction and backscatter lidar data by inversion with regularization: theory. *Applied Optics*, 38(12):2346–2357, doi:10.1364/AO.38.002346.
- Müller, J. W. (1973). Dead-time problems. *Nuclear Instruments and Methods*, 112(1-2):47–57, doi:10.1016/0029-554x(73)90773-8.
- National Oceanic and Atmospheric Administration (NOAA) (1976). U.S. Standard Atmosphere. Technical Report NOAA Doc. S/T 76-1562, U.S. Government Printing Office, Washington, D.C.
- Naud, C. M., Müller, J. P., and Clothiaux, E. E. (2003). Comparison between active sensor and radiosonde cloud boundaries over the ARM Southern Great Plains site. *Journal of Geophysical Research — Atmospheres*, 108:4140, doi:10.1029/2002jd002887.
- Network for the Detection of Atmospheric Composition Change (NDACC) [Rosen, J.] (1989 – 2000). Backscattersonde data (Laramie, WY; Lauder, New Zealand; Thule, Greenland). Supported by NSF grants OPP-9423285, ATM-9500186.

- Newsom, R. K., Turner, D. D., Mielke, B., Clayton, M., Ferrare, R., and Sivaraman, C. (2009). Simultaneous analog and photon counting detection for Raman lidar. *Applied Optics*, 48(20):3903–3914, doi:10.1364/AO.48.003903.
- O'Dowd, C., Ceburnis, D., Ovadnevaite, J., Martucci, G., Bialek, J., Monahan, C., Berresheim, H., Vaishya, A., Grigas, T., Jennings, S. G., McVeigh, P., Varghese, S., Flanagan, R., Martin, D., Moran, E., Lambkin, K., Semmler, T., Perrino, C., and McGrath, R. (2012). The Eyjafjallajökull ash plume — Part I: Physical, chemical and optical characteristics. *Atmospheric Environment*, 48:129–142, doi:10.1016/j.atmosenv.2011.07.004.
- Oke, T. R. (1987). *Boundary Layer Climates*. Cambridge University Press, London, UK, second edition.
- Pahlow, M., Feingold, G., Jefferson, A., Andrews, E., Ogren, J. A., Wang, J., Lee, Y. N., Ferrare, R. A., and Turner, D. D. (2006). Comparison between lidar and nephelometer measurements of aerosol hygroscopicity at the Southern Great Plains Atmospheric Radiation Measurement site. *Journal of Geophysical Research — Atmospheres*, 111:D05S15, doi:10.1029/2004jd005646.
- Palmer, P. I., Barnett, J. J., Eyre, J. R., and Healy, S. B. (2000). A nonlinear optimal estimation inverse method for radio occultation measurements of temperature, humidity, and surface pressure. *Journal of Geophysical Research — Atmospheres*, 105(D13):17513–17526, doi:10.1029/2000jd900151.
- Pappalardo, G., Amodeo, A., Mona, L., and Pandolfi, M. (2005). Systematic measurements of the aerosol extinction-to-backscatter ratio. During *Lidar Remote Sensing for Industry and Environmental Monitoring V*, editors: Singh, U. N. and Mizutani, K., volume 5653, pages 77–87, Bellingham, WA. SPIE.
- Pornsawad, P., D'Amico, G., Böckmann, C., Amodeo, A., and Pappalardo, G. (2012). Retrieval of aerosol extinction coefficient profiles from Raman lidar data by inversion method. *Applied Optics*, 51(12):2035–2044, doi:10.1364/AO.51.002035.
- Pounder, N. L., Hogan, R. J., Várnai, T., Battaglia, A., and Cahalan, R. F. (2012). A variational method to retrieve the extinction profile in liquid clouds using multiple-field-of-view lidar. *Journal of Applied Meteorology and Climatology*, 51(2):350–365, doi:10.1175/jamc-d-10-05007.1.

- Povey, A. C., Grainger, R. G., Peters, D. M., Agnew, J. L., and Rees, D. (2012). Estimation of a lidar's overlap function and its calibration by nonlinear regression. *Applied Optics*, 51(21):5130–5143, doi:10.1364/AO.51.005130.
- Press, W. H., Vetterling, W. T., Teukolsky, S. A., and Flannery, B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New Delhi, India, second edition.
- Reagan, J. A., McCormick, M. P., and Spinhirne, J. D. (1989). Lidar sensing of aerosols and clouds in the troposphere and stratosphere. *Proceedings of the IEEE*, 77(3):433–448, doi:10.1109/5.24129.
- Rodgers, C. D. (2000). *Inverse Methods for Atmospheric Sounding: Theory and Practice*, volume 2 of *Series on Atmospheric, Oceanic, and Planetary Physics*. World Scientific, Singapore, second edition.
- Rosen, J., Young, S., Laby, J., Kjome, N., and Gras, J. (2000). Springtime aerosol layers in the free troposphere over Australia: Mildura Aerosol Tropospheric Experiment (MATE 98). *Journal of Geophysical Research — Atmospheres*, 105(D14):17833–17842, doi:10.1029/2000jd900208.
- Rosen, J. M. and Kjome, N. T. (1991). Backscattersonde — a new instrument for atmospheric aerosol research. *Applied Optics*, 30(12):1552–1561, doi:10.1364/AO.30.001552.
- Sasano, Y., Shimizu, H., Takeuchi, N., and Okuda, M. (1979). Geometrical form-factor in the laser-radar equation — experimental determination. *Applied Optics*, 18(23):3908–3910, doi:10.1364/ao.18.003908.
- Schmid, B., Ferrare, R., Flynn, C., Elleman, R., Covert, D., Strawa, A., Welton, E., Turner, D., Jonsson, H., Redemann, J., Eilers, J., Ricci, K., Hallar, A. G., Clayton, M., Michalsky, J., Smirnov, A., Holben, B., and Barnard, J. (2006). How well do state-of-the-art techniques measuring the vertical profile of tropospheric aerosol extinction compare? *Journal of Geophysical Research — Atmospheres*, 111:D05S07, doi:10.1029/2005jd005837.
- Schumann, U., Weinzierl, B., Reitebuch, O., Schlager, H., Minikin, A., Forster, C., Baumann, R., Sailer, T., Graf, K., Mannstein, H., Voigt, C., Rahm, S., Simmet, R., Scheibe, M., Lichtenstern, M., Stock, P., Rüba, H., Schäuble, D., Tafferner, A., Rautenhaus, M., Gerz, T., Ziereis, H., Krautstrunk, M., Mallaun, C., Gayet, J. F.,

- Lieke, K., Kandler, K., Ebert, M., Weinbruch, S., Stohl, A., Gasteiger, J., Gross, S., Freudenthaler, V., Wiegner, M., Ansmann, A., Tesche, M., Olafsson, H., and Sturm, K. (2011). Airborne observations of the Eyjafjalla volcano ash cloud over Europe during air space closure in April and May 2010. *Atmospheric Chemistry and Physics*, 11(5):2245–2279, doi:10.5194/acp-11-2245-2011.
- Science and Technology Facilities Council (STFC) [Wrench, C.L.] (2006 – 2011). Chilbolton Facility for Atmospheric and Radio Research (CFARR) data. Available from NCAS British Atmospheric Data Centre at http://badc.nerc.ac.uk/view/badc.nerc.ac.uk___ATOM__dataent_chobs (last accessed 9 Nov 2010).
- Seibert, P., Beyrich, F., Gryning, S. E., Joffre, S., Rasmussen, A., and Tercier, P. (2000). Review and intercomparison of operational methods for the determination of the mixing height. *Atmospheric Environment*, 34(7):1001–1027, doi:10.1016/s1352-2310(99)00349-0.
- Seinfeld, J. and Pandis, S. (1998). *Atmospheric Chemistry and Physics*. Wiley, New York, first edition.
- Shcherbakov, V. (2007). Regularized algorithm for Raman lidar data processing. *Applied Optics*, 46(22):4879–4889, doi:10.1364/AO.46.004879.
- She, C. Y., Latifi, H., Yu, J. R., Alvarez, R. J., Bills, R. E., and Gardner, C. S. (1990). 2-frequency lidar technique for mesospheric Na temperature-measurements. *Geophysical Research Letters*, 17(7):929–932, doi:10.1029/GL017i007p00929.
- Stelmaszczyk, K., Dell’Aglio, M., Chudzynski, S., Stacewicz, T., and Wöste, L. (2005). Analytical function for lidar geometrical compression form-factor calculations. *Applied Optics*, 44(7):1323–1331, doi:10.1364/ao.44.001323.
- Steyn, D. G., Baldi, M., and Hoff, R. M. (1999). The detection of mixed layer depth and entrainment zone thickness from lidar backscatter profiles. *Journal of Atmospheric and Oceanic Technology*, 16(7):953–959, doi:10.1175/1520-0426(1999)016<0953:TDOMLD>2.0.CO;2.
- Stull, R. (1988). *An Introduction to Boundary Layer Meteorology*. Kluwer Academic, London.

- Sugimoto, N., Matsui, I., Shimizu, A., Nishizawa, T., Hara, Y., Chenbo, X., Uno, I., Yumimoto, K., Zifa, W., and Soon-Chang, Y. (2008). Lidar network observations of tropospheric aerosols. *Proceedings of the SPIE*, 7153:71530A, doi:10.1117/12.806540.
- Tomine, K., Hirayama, C., Michimoto, K., and Takeuchi, N. (1989). Experimental determination of the crossover function in the laser-radar equation for days with a light mist. *Applied Optics*, 28(12):2194–2195, doi:10.1364/AO.28.002194.
- Turner, D. D., Ferrare, R. A., Brasseur, L. A. H., and Feltz, W. F. (2002). Automated retrievals of water vapor and aerosol profiles from an operational Raman lidar. *Journal of Atmospheric and Oceanic Technology*, 19(1):37–50, doi:10.1175/1520-0426(2002)019<0037:AROWVA>2.0.CO;2.
- UK Meteorological Office [Parton, G.] (2006 – 2011). Met Office Global Radiosonde Data. Available from the NCAS British Atmospheric Data Centre at http://badc.nerc.ac.uk/view/badc.nerc.ac.uk__ATOM__dataent_GLOBRADS (last accessed 9 Feb 2011).
- Vaughan, M., Young, S., Winker, D., Powell, K., Omar, A., Liu, Z. Y., Hu, Y. X., and Hostetler, C. (2004). Fully automated analysis of space-based lidar data: an overview of the CALIPSO retrieval algorithms and data products. *Laser Radar Techniques for Atmospheric Sensing*, 5575:16–30, doi:10.1117/12.572024.
- Velotta, R., Bartoli, B., Capobianco, R., Fiorani, L., and Spinelli, N. (1998). Analysis of the receiver response in lidar measurements. *Applied Optics*, 37(30):6999–7007, doi:10.1364/AO.37.006999.
- Veselovskii, I., Kolgotin, A., Griaznov, V., Müller, D., Wandinger, U., and Whiteman, D. N. (2002). Inversion with regularization for the retrieval of tropospheric aerosol parameters from multiwavelength lidar sounding. *Applied Optics*, 41(18):3685–3699, doi:10.1364/ao.41.003685.
- von Zahn, U., von Cossart, G., Fiedler, J., Fricke, K. H., Nelke, G., Baumgarten, G., Rees, D., Hauchecorne, A., and Adolfsen, K. (2000). The ALOMAR Rayleigh/Mie/Raman lidar: objectives, configuration, and performance. *Annales Geophysicae*, 18(7):815–833, doi:10.1007/s005850000210.

- Wandinger, U. and Ansmann, A. (2002). Experimental determination of the lidar overlap profile with Raman lidar. *Applied Optics*, 41(3):511–514, doi:10.1364/AO.41.000511.
- Wandinger, U., Ansmann, A., Reichardt, J., and Deshler, T. (1995). Determination of stratospheric aerosol microphysical properties from independent extinction and backscattering measurements with a Raman lidar. *Applied Optics*, 34(36):8315–8329, doi:10.1364/AO.34.008315.
- Watts, P. D., Bennartz, R., and Fell, F. (2011). Retrieval of two-layer cloud properties from multispectral observations using optimal estimation. *Journal of Geophysical Research — Atmospheres*, 116:D16203, doi:10.1029/2011jd015883.
- Weitkamp, C. (2005). *Lidar: Range-Resolved Optical Remote Sensing of the Atmosphere*. Springer Series in Optical Sciences. Springer, New York.
- Welton, E. J., Campbell, J. R., Spinhirne, J. D., and Scott, V. S. (2001). Global monitoring of clouds and aerosols using a network of micro-pulse lidar systems. During *Conference on Lidar Remote Sensing for Industry and Environment Monitoring*, volume 4153, pages 151–158.
- Whiteman, D. N. (1999). Application of statistical methods to the determination of slope in lidar data. *Applied Optics*, 38(15):3360–3369, doi:10.1364/ao.38.003360.
- Whiteman, D. N. (2003). Examination of the traditional Raman lidar technique. I. Evaluating the temperature-dependent lidar equations. *Applied Optics*, 42(15):2571–2592, doi:10.1364/AO.42.002571.
- Whiteman, D. N., Melfi, S. H., and Ferrare, R. A. (1992). Raman lidar system for the measurement of water-vapor and aerosols in the Earth’s atmosphere. *Applied Optics*, 31(16):3068–3082, doi:10.1364/AO.31.003068.
- Wiegner, M., Emeis, S., Freudenthaler, V., Heese, B., Junkermann, W., Munkel, C., Schäfer, K., Seefeldner, M., and Vogt, S. (2006). Mixing layer height over Munich, Germany: Variability and comparisons of different methodologies. *Journal of Geophysical Research — Atmospheres*, 111(D13):D13201, doi:10.1029/2005jd006593.
- Woodhouse, I. and Agnew, J. (2010 – 2011). Chilbolton AERONET Level 2.0 real time data. Available from NASA GSFC at <http://aeronet.gsfc.nasa.gov> (last accessed 31 Mar 2013).

- Young, S. A. (1995). Analysis of lidar backscatter profiles in optically thin clouds. *Applied Optics*, 34(30):7019–7031, doi:10.1364/AO.34.007019. doi:10.1364/AO.34.007019.
- Zeyn, J., Lahmann, W., and Weitkamp, C. (1996). Remote daytime measurements of tropospheric temperature profiles with a rotational Raman lidar. *Optics Letters*, 21(16):1301–1303, doi:10.1364/ol.21.001301.