Chapter 4

Adjoint-state methods

As explained in section (3.4), the adjoint F^* of the linearized forward (modeling) operator F plays an important role in the formula of the functional gradient $\frac{\delta J}{\delta m}$ of the least-squares cost function J:

$$\frac{\delta J}{\delta m}[m] = F^*(\mathcal{F}[m] - d).$$

While F is the basic linear map from model space to data space, F^* is the basic linear map from data space to model space. F^* is not only the building block of iterative optimization schemes, but the mere application of F^* to data is the simplest form of "imaging". For instance, when the initial guess $m^{(0)} = m_0$ is a smooth background model reasonably close to the true solution m, and when there is a sufficiently large number of receivers and/or sources, the first iteration of gradient descent,

$$m^{(1)} = \alpha F^*(d - \mathcal{F}[m_0]),$$

often gives a good "image" of the scatterers (somewhat close to the actual εm_1). For this reason, F^* is often called the *imaging operator*.

It should also be noted that F^* behaves not entirely unlike F^{-1}, i.e., F is somewhat close to being unitary. This statement does not have a rigorous formulation of the form $\|F^*F - I\| \leq (\ldots)$, but rather of the form "F^*F does not move singularities around like F or F^* do". More details on the microlocal aspects of this question will be given in chapter 8.1.

Forming the full matrix $F = \frac{\delta \mathcal{F}}{\delta m}$ and transposing it is not a practical way to compute F^*. The adjoint-state method provides an elegant solution to this problem, resulting in what is called the "imaging condition".

4.1 The imaging condition

For any $d_r(t)$ function of the receiver index r and time t, and $m(x)$ function of position x (here m and d are two arbitrary functions, not necessarily linked to one another by the forward model), we have

$$\langle d, Fm \rangle = \langle F^* d, m \rangle.$$

The inner product on the left is in data space,

$$\langle d, Fm \rangle = \sum_r \int_0^T d_r(t) u(x_r, t)\, dt, \qquad u = Fm,$$

while the inner product on the right is in model space.

$$\langle F^* d, m \rangle = \int_{\mathbb{R}^n} (F^* d)(x) m(x)\, dx.$$

The relation $u = Fm$ is implicitly encoded by the two equations

$$\left(m_0 \frac{\partial^2}{\partial t^2} - \Delta \right) u = -m \frac{\partial^2 u_0}{\partial t^2},$$

$$\left(m_0 \frac{\partial^2}{\partial t^2} - \Delta \right) u_0 = f.$$

Note that a single right-hand side generates u_0, and that we have omitted the source subscript s in this section; we will return to multiples sources shortly.

The argument that isolates and makes explicit the contribution of m in $\sum_r \int_0^T d_r(t) u(x_r, t)\, dt$ is one of integration by parts. In order to integrate by parts in x, we need to turn the sum over receivers into an integral. This can be achieved by considering a distributional extended dataset where each measurement $d_r(t)$ is accompanied by a Dirac delta located at x_r:

$$d_{\text{ext}}(x, t) = \sum_r d_r(t) \delta(x - x_r).$$

We then have

$$\langle d, Fm \rangle = \int_{\mathbb{R}^n} \int_0^T d_{\text{ext}}(x, t) u(x, t)\, dx dt.$$

In order to use the wave equation for u, a copy of the differential operator $\left(m_0 \frac{\partial^2}{\partial t^2} - \Delta \right)$ needs to materialize. This is done by considering an auxiliary field $q(x,t)$ that solves the same wave equation with $d_{\text{ext}}(x,t)$ as a right-hand side:

$$\left(m_0 \frac{\partial^2}{\partial t^2} - \Delta \right) q(x,t) = d_{\text{ext}}(x,t), \qquad x \in \mathbb{R}^n, \tag{4.1}$$

with as-yet unspecified "boundary conditions" in time. Substituting this expression for $d_{\text{ext}}(x,t)$, and integrating by parts both in space and in time reveals

$$\begin{aligned} \langle d, Fm \rangle = & \int_V \int_0^T q(x,t) \left(m_0 \frac{\partial^2}{\partial t^2} - \Delta \right) u(x,t)\, dxdt \\ & + \int_V m_0 \frac{\partial q}{\partial t} u|_0^T \, dx - \int_V m_0 q \frac{\partial u}{\partial t}|_0^T \, dx \\ & + \int_{\partial V} \int_0^T \frac{\partial q}{\partial n} u \, dS_x dt - \int_{\partial V} \int_0^T q \frac{\partial u}{\partial n} \, dS_x dt, \end{aligned}$$

where V is a volume that extends to the whole of $\mathbb{R}^n$, and ∂V is the boundary of V — the equality then holds in the limit of $V = \mathbb{R}^n$.

The boundary terms over ∂V vanish in the limit of large V by virtue of the fact that they involve u – a wavefield created by localized functions f, m, u_0 and which does not have time to travel arbitrarily far within a time $[0, T]$. The boundary terms at $t = 0$ vanish due to $u|_{t=0} = \frac{\partial u}{\partial t}|_{t=0} = 0$. As for the boundary terms at $t = T$, they only vanish if we impose

$$q|_{t=T} = \frac{\partial q}{\partial t}|_{t=T} = 0. \tag{4.2}$$

Since we are only interested in the values of $q(x,t)$ for $0 \le t \le T$, the above are *final conditions* rather than initial conditions, and the equation (4.1) is run *backward in time*. The wavefield q is called *adjoint* field, or adjoint state. The equation (4.1) is itself called adjoint equation. Note that q is *not* in general the physical field run backward in time (because of the limited sampling at the receivers), instead, it is introduced purely out of computational convenience.

The analysis needs to be modified when boundary conditions are present. Typically, homogeneous Dirichlet and Neumann boundary conditions should the same for u_0 and for q – a choice which will manifestly allow to cancel out the boundary terms in the reasoning above – but absorbing boundary conditions involving time derivatives need to be properly time-reversed as well. A systematic, painless way of solving the adjoint problem is to follow the following sequence of three steps: (i) time-reverse the data d_{ext} at each receiver, (ii) solve the wave equation forward in time with this new right-hand-side, and (iii) time-reverse the result at each point x.

We can now return to the simplification of the left-hand-side,

$$\begin{aligned}\langle d, Fm\rangle &= \int_{\mathbb{R}^n}\int_0^T q(x,t)\left(m_0\frac{\partial^2}{\partial t^2} - \Delta\right)u(x,t)\,dxdt\\ &= -\int_{\mathbb{R}^n}\int_0^T q(x,t)m(x)\frac{\partial^2 u_0}{\partial t^2}\,dxdt\end{aligned}$$

This quantity is also supposed to be $\langle m, F^*d\rangle$, regardless of m, so we conclude

$$(F^*d)(x) = -\int_0^T q(x,t)\frac{\partial^2 u_0}{\partial t^2}\,dt. \tag{4.3}$$

This equation is called the *imaging condition*: it expresses the action of F^* on d as the succession of the following steps:

1. Place data $d_r(t)$ at the location of the receivers with point masses to get d_{ext};

2. Use d_{ext} as the right-hand side in the adjoint wave equation to get the adjoint, backward field q;

3. Simulate the incident, forward field u_0; and finally

4. Take the time integral of the product of the forward field u_0 (differentiated twice in t), and the backward field q, for each x independently.

The result is a function of x which sometimes serves the purpose of image, and may sometimes be called $I_m(x)$. Note that we have not performed a full inversion; if d are measured data, then I_m is not the model m that gave rise to d. In seismology, the imaging condition (4.3) is called reverse-time migration, or simply *migration*. In radar, the imaging condition does not have a particular name, but in the next chapter we will encounter a simplification of (4.3) called backprojection.

If we now restore the presence of multiple sources, the wavefields u, u_0, and u_1 will depend on the source index s. The source term f_s — typically of the form $w(t)\delta(x - x_s)$ — is in the right-hand side of the wave equations for

u_0 and u, while u_1 implicitly depends on f_s through u_0. For a fixed source s, we denote

$$u_s = \mathcal{F}_s[m], \qquad u_{0,s} = \mathcal{F}_s[m_0], \qquad u_{1,s} = F_s m_1,$$

while we continue to denote $u = \mathcal{F}[m]$, $u_0 = \mathcal{F}[m_0]$ and $u_1 = F m_1$ for the *collection* of such wavefields over s.

The data inner-product now has an additional sum over s, namely

$$\langle d, Fm \rangle = \sum_s \sum_r \int_0^T d_{r,s}(t) u_s(x_r, t)\, dt.$$

The formula for F^* can be obtained by taking adjoints one s at a time, namely

$$\begin{aligned}\langle F^* d, m \rangle = \langle d, Fm \rangle &= \sum_s \langle d_s, F_s m \rangle \\ &= \sum_s \langle F_s^* d_s, m \rangle \\ &= \langle \sum_s F_s^* d_s, m \rangle,\end{aligned}$$

hence

$$F^* = \sum_s F_s^*.$$

More explicitly, in terms of the imaging condition,

$$(F^* d)(x) = -\sum_s \int_0^T q_s(x,t) \frac{\partial^2 u_{0,s}}{\partial t^2}(x,t)\, dt, \tag{4.4}$$

where the adjoint field q_s is relative to the source s:

$$\left(m_0 \frac{\partial^2}{\partial t^2} - \Delta \right) q_s(x,t) = d_{\mathrm{ext},s}(x,t).$$

The sum over s in the new imaging condition (4.4) is sometimes called a *stack*. It is often the case that particular images $F_s^* d$ are not very informative on their own, but a stack uses the redundancy in the data to bring out the information and reveal more details.

The mathematical tidbit underlying stacks is that the operation of creating a vector $(x, x, \ldots, x)$ out of a single number x has for adjoint the operation of summing the components of a vector.

4.2 The imaging condition in the frequency domain

We now modify the exposition to express both the adjoint-state equation and the imaging condition in the frequency (ω) domain. The nugget in this section is that complex conjugation in ω corresponds to time reversal. We assume a single source for simplicity.

We are again interested in finding F^* such that $\langle d, Fm\rangle = \langle F^* d, m\rangle$ for all generic d and m. The data inner product $\langle d, Fm\rangle$ can be expressed in the frequency domain by means of the Parseval formula,

$$\langle d, Fm\rangle = 2\pi \sum_r \int_{\mathbb{R}} \widehat{d_r}(\omega) \overline{\widehat{(Fm)}(x_r,\omega)} d\omega = \sum_r \int d_r(t)(Fm)(x_r,t)dt.$$

The complex conjugate is important, now that we are in the frequency domain – though since the overall quantity is real, it does not matter which of the two integrand's factors it is placed on. As previously, we pass to the extended dataset

$$\widehat{d_{\mathrm{ext}}}(x,\omega) = \sum_r \widehat{d_r}(\omega)\delta(x - x_r),$$

and turn the sum over r into an integral over x. The linearized scattered field is

$$\widehat{(Fm)}(x_r,\omega) = \int \widehat{G}(x,y;\omega) m(y) \omega^2 \widehat{u_0}(y,\omega)\, dy. \tag{4.5}$$

To simplify the resulting expression of $\langle d, Fm\rangle$, we let

$$\widehat{q}(x,\omega) = \int \overline{\widehat{G}(y,x;\omega)} \widehat{d_{\mathrm{ext}}}(y,\omega)\, dy. \tag{4.6}$$

(Note that Green's functions are always symmetric under the swap of x and y, as we saw in a special case in one of the exercises in chapter 1.) It follows that

$$\langle d, Fm\rangle = \int m(y) \left[2\pi \int_{\mathbb{R}} \widehat{q}(y,\omega) \omega^2 \overline{\widehat{u_0}(y,\omega)}\, d\omega \right] dy,$$

hence

$$F^* d(y) = 2\pi \int_{\mathbb{R}} \widehat{q}(y,\omega) \omega^2 \overline{\widehat{u_0}(y,\omega)}\, d\omega. \tag{4.7}$$

This equation is the same as (4.3), by Parseval's identity. Equation (4.6) is the integral version of (4.1) in the frequency domain. The complex conjugation of $\widehat{G}$ in (4.6) is the expression in the frequency domain of the fact that the adjoint equation is solved backwards in time[1]. We can alternatively interpret $\widehat{q} = \overline{\widehat{G}}\widehat{d_{\text{ext}}}$ by applying an extra conjugate, namely $\overline{\widehat{q}} = \widehat{G}\overline{\widehat{d_{\text{ext}}}}$, which can be read as the sequence of operations: (i) time-reverse d_{ext}, (ii) propagate it forward in time, and (iii) time-reverse the result. This is the same prescription as in the time-domain case, and offers the added advantage of not having to rethink the boundary condition for the backward equation.

The integral in t in (4.3) is over $[0, T]$ because such is the support of $q \frac{\partial^2 u_0}{\partial t^2}$. The integral in ω in (4.7) is over $\mathbb{R}$. It is tempting to truncate this integral to "the frequencies that have been measured" — but that is strictly speaking incompatible with the limits on t (for the same reason that a function compactly supported in time cannot also be compactly supported in frequency.) Careful consideration of cutoffs is needed to control the accuracy of a truncation in ω.

Equation (4.7) is valuable for a few different reasons:

- It can be further simplified to highlight its geometrical content as an approximate curvilinear integral, such as explained in the next chapter;

- The integral over ω can be deliberately restricted in the scope of descent iterations, so as to create sweeps over frequencies. This is sometimes important to deal with the lack of convexity of full inversion; see chapter 9.

4.3 The general adjoint-state method

In this section we explain how to use the adjoint-state method to compute the first and second variations of an objective function $J[u(m)]$ in a parameter m, when u is constrained by the equation $L(m)u = f$, where $L(m)$ is a linear operator that depends on m.

[1]Time reversal of any real-valued function becomes conjugation in the Fourier domain:

$$\overline{\widehat{f}(\omega)} = \overline{\int e^{i\omega t} f(t)\, dt} = \int e^{-i\omega t} f(t) dt = \int e^{i\omega t} f(-t)\, dt.$$

We refer to "u space", "m space", and "f space" for the respective L^2 spaces containing u, m, and f. The first variation of J is simply

$$\frac{\delta J}{\delta m} = \langle \frac{\delta J}{\delta u}, \frac{\delta u}{\delta m} \rangle_u, \tag{4.8}$$

where the inner product pairs δu in each equation, hence acts in u space. Notice that $\delta u / \delta m$ is an operator acting in m space and returning a function in u space[2].

If we were interested in computing the directional derivative of J in some direction m, namely $\langle \frac{\delta J}{\delta m}, m \rangle$, then we would simply swap the m-space inner product with the u-space inner product from (4.8), and recognize that $u = \frac{\delta u}{\delta m} m$ is easily accessible by solving the linearized version of the equation $L(m)u = f$. This result is straightforward to obtain, and the adjoint-state method is not necessary.

The interesting problem is that of computing $\frac{\delta J}{\delta m}$ as a function in m-space. In principle, equation (4.8) is all that is needed for this purpose, except that explicitly computing the full kernel of the operator $\frac{\delta u}{\delta m}$ can be highly inefficient in comparison to the complexity of specifying the function $\frac{\delta J}{\delta m}$.

The adjoint-state method is a very good way of eliminating $\frac{\delta u}{\delta m}$ so that $\frac{\delta J}{\delta m}$ can be computed in more favorable complexity. In order to achieve this, differentiate the "state equation" $L(m)u = f$ with respect to m to get

$$\frac{\delta L}{\delta m} u + L \frac{\delta u}{\delta m} = 0.$$

We see that $\frac{\delta u}{\delta m}$ can be eliminated by composition with L on the left. The main idea of the adjoint-state method is that a copy of L can materialize in (4.8) provided the other factor, $\frac{\delta J}{\delta u}$, is seen as the *adjoint* of L applied to some field q,

$$L^* q = \frac{\delta J}{\delta u}, \qquad \text{(adjoint-state equation)} \tag{4.9}$$

with q naturally called the adjoint field. Then,

$$\begin{aligned} \frac{\delta J}{\delta m} &= \langle L^* q, \frac{\delta u}{\delta m} \rangle_u = \langle q, L \frac{\delta u}{\delta m} \rangle_f \\ &= -\langle q, \frac{\delta L}{\delta m} u \rangle_f \qquad \text{(imaging condition)} \end{aligned} \tag{4.10}$$

[2]It relates to what we called $F = \delta \mathcal{F} / \delta m$ earlier by the operator S of sampling at the receivers, via $\mathcal{F} = S u$ or $\delta \mathcal{F} / \delta m = S \delta u / \delta m$.

This latter expression is often much easier to compute than (4.8).

Example 1. *In the setting of section 4.1, m is a function of x; u and f are functions of (x,t). The state equation $L(m)u = f$ is the forward wave equation $m\partial_t^2 u - \Delta u = f$ with zero initial conditions. When we evaluate all our quantites at m_0, then u becomes u_0, the incident field. The adjoint-state equation $L^* q = \frac{\delta J}{\delta u}$ is the backward wave equation $m\partial_t^2 q - \Delta q = \frac{\delta J}{\delta u}$ with zero final conditions. The least-squares cost function is $J[u] = \frac{1}{2}\|Su - d\|_2^2$ with S the operator of sampling at the receivers, so that the adjoint source $\frac{\delta J}{\delta u} = S^*(Su - d)$ is the data residual extended to a function of x and t.*

The quantity $\frac{\delta L}{\delta m} u$ is a multiplication operator from m-space to f-space (which takes the function m to the function $m\partial_t^2 u$), expressed in coordinates as

$$(\frac{\delta L}{\delta m(y)} u)(x,t) = \delta(x-y)\, \partial_t^2 u(x,t), \qquad (x \in R^3, y \in \mathbb{R}^3.)$$

Using the formula above, $-\langle q, \frac{\delta L}{\delta m} u\rangle_f$ becomes the usual imaging condition $-\int q(x,t)\partial_t^2 u(x,t)\, dt$.

The adjoint-state method also allows to compute second variations. We readily compute the functional Hessian of J as

$$\frac{\delta^2 J}{\delta m \delta m'} = \langle \frac{\delta u}{\delta m}, \frac{\delta^2 J}{\delta u \delta u'} \frac{\delta u'}{\delta m'} \rangle_u + \langle \frac{\delta J}{\delta u}, \frac{\delta^2 u}{\delta m \delta m'} \rangle_u,$$

where the u-space inner product pairs the δu factors in each expression.

This Hessian is an object with two free variables[3] m and m'. If we wish to view it as a bilinear form, and compute its action $\langle m, \frac{\delta^2 J}{\delta m \delta m'} m' \rangle_m$ on two functions m and m', then it suffices to pass those functions inside the u-space inner product to get

$$\langle u, \frac{\delta^2 J}{\delta u \delta u'} u' \rangle_u + \langle \frac{\delta J}{\delta u}, v \rangle_u.$$

The three functions $u = \frac{\delta u}{\delta m} m$, $u' = \frac{\delta u'}{\delta m'} m'$, and $v = \langle m, \frac{\delta^2 u}{\delta m \delta m'} m' \rangle$ can be computed by solving simple (linearized) equations derived from the state equation $L(m)u = f$. (See a few exercises at the end of this chapter.) The

[3]Here m and m' are functions, hence can be seen as independent variables, i.e., points, in a function space. Free, or unpaired variables are to functional calculus what free indices are to vector calculus: they are "open slots" available to be filled by pairing with a function/vector in an inner product.

adjoint-state method is not needed for evaluating the Hessian as a bilinear form.

On the other hand, we are generally not interested in computing the full matrix representation of $\frac{\delta^2 J}{\delta m \delta m'}$ with row index m and column index m' either: this object is too large to store.

The most interesting question is the computation of the action of the Hessian as an operator on a function, say m'. While m' is paired, the variable m remains free, hence the result is a function in m-space. This problem can be solved by the second-order adjoint-state method, which we now explain.

As earlier, we can pass m' inside the u-space inner product to get

$$\frac{\delta^2 J}{\delta m \delta m'} m' = \langle \frac{\delta u}{\delta m}, \frac{\delta^2 J}{\delta u \delta u'} u' \rangle + \langle \frac{\delta J}{\delta u}, \frac{\delta^2 u}{\delta m \delta m'} m' \rangle, \tag{4.11}$$

with $u' = \frac{\delta u'}{\delta m'} m'$ easily computed. However, two quantities are not immediately accessible:

1. the remaining $\frac{\delta u}{\delta m}$ factor with the un-paired m variable, and

2. the second-order $\frac{\delta^2 u}{\delta m \delta m'} m'$ factor.

Both quantities have two free variables u and m, hence are too large to store, let alone compute.

The second-order factor can be handled as follows. The second variation of the equation $L(m)u = f$ is

$$\frac{\delta^2 L}{\delta m \delta m'} u + \frac{\delta L}{\delta m} \frac{\delta u}{\delta m'} + \frac{\delta L}{\delta m'} \frac{\delta u}{\delta m} + L \frac{\delta^2 u}{\delta m \delta m'} = 0.$$

We see that the factor $\frac{\delta^2 u}{\delta m \delta m'}$ can be eliminated provided L is applied to it on the left. We follow the same prescription as in the first-order case, and define a *first adjoint field* q_1 such that

$$L^* q_1 = \frac{\delta J}{\delta u}. \qquad \text{(1st adjoint-state equation)} \tag{4.12}$$

A substitution in (4.11) reveals

$$\begin{aligned} \frac{\delta^2 J}{\delta m \delta m'} m' = &\langle \frac{\delta u}{\delta m}, \frac{\delta^2 J}{\delta u \delta u'} u' \rangle_u - \langle q_1, \left(\frac{\delta L}{\delta m'} m' \right) \frac{\delta u}{\delta m} \rangle_f \\ &- \langle q_1, \left(\frac{\delta^2 L}{\delta m \delta m'} m' \right) u + \frac{\delta L}{\delta m} u' \rangle_f, \end{aligned}$$

with $u' = \frac{\delta u}{\delta m'} m'$, as earlier. The term on the last row can be computed as is; all the quantities it involves are accessible. The two terms in the right-hand-side of the first row can be rewritten to isolate $\delta u / \delta m$, as

$$\langle \frac{\delta^2 J}{\delta u \delta u'} u' - \left(\frac{\delta L}{\delta m'} m' \right)^* q_1, \, \frac{\delta u}{\delta m} \rangle_u.$$

In order to eliminate $\frac{\delta u}{\delta m}$ by composing it with L on the left, we are led to defining a *second adjoint-state field* q_2 via

$$L^* q_2 = \frac{\delta^2 J}{\delta u \delta u'} u' - \left(\frac{\delta L}{\delta m'} m' \right)^* q_1. \qquad \text{(2nd adjoint-state equation)} \quad (4.13)$$

All the quantities in the right-hand side are available. It follows that

$$\langle L^* q_2, \frac{\delta u}{\delta m} \rangle_u = \langle q_2, L \frac{\delta u}{\delta m} \rangle = - \langle q_2, \frac{\delta L}{\delta m} u \rangle_f.$$

Gathering all the terms, we have

$$\frac{\delta^2 J}{\delta m \delta m'} m' = - \langle q_2, \frac{\delta L}{\delta m} u \rangle_f - \langle q_1, \left(\frac{\delta^2 L}{\delta m \delta m'} m' \right) u + \frac{\delta L}{\delta m} u' \rangle_f,$$

with q_1 obeying (4.12) and q_2 obeying (4.13).

Example 2. *In the setting of section 4.1, call the base model $m = m_0$ and the model perturbation $m' = m_1$, so that $u' = \frac{\delta u}{\delta m}[m_0] m'$ is the solution u_1 of the linearized wave equation $m_0 \partial_t^2 u_1 - \Delta u_1 = -m_1 \partial_t^2 u_0$ with $u_0 = L(m_0) u_0 = f$ as previously. The first variation $\frac{\delta L}{\delta m}$ is the same as explained earlier, while the second variation $\frac{\delta^2 L}{\delta m \delta m'}$ vanishes since the wave equation is linear in m. Thus if we let $H = \frac{\delta^2 J}{\delta m \delta m'}$ for the Hessian of J as an operator in m-space, we get*

$$H m_1 = - \int q_2(x,t) \partial_t^2 u_0(x,t) \, dt - \int q_1(x,t) \partial_t^2 u_1(x,t) \, dt.$$

The first adjoint field is the same as in the previous example, namely q_1 solves the backward wave equation

$$m_0 \partial_t^2 q_1 - \Delta q_1 = S^*(S u_0 - d), \qquad \textit{zero f.c.}$$

To get the equation for q_2, notice that $\frac{\delta^2 J}{\delta u \delta u'} = S^* S$ and $\left(\frac{\delta L}{\delta m'} m_1 \right)^* = \frac{\delta L}{\delta m'} m_1 = m_1 \partial_t^2$. Hence q_2 solves the backward wave equation

$$m_0 \partial_t^2 q_2 - \Delta q_2 = S^* S u_1 - m_1 \partial_t^2 q_1, \qquad \textit{zero f.c.}$$

Note that the formulas (3.14) (3.15) for H that we derived in an exercise in the previous chapter are still valid, but they do not directly allow the computation of H as an operator acting in m-space. The approximation $H \simeq F^*F$ obtained by neglecting the last term in (3.14) is recovered in the context of second-order adjoints by letting $q_1 = 0$.

4.4 The adjoint state as a Lagrange multiplier

The adjoint field q was introduced in a somewhat opportunistic and artificial way in earlier sections. In this section, we show that it has the interpretation of a Lagrange multiplier in a constrained optimization framework, where the wave equation serves as the constraint. The problem is the same as in the previous section, namely to compute the functional gradient of $J[u(m)]$ with respect to m – and the resulting expression is the same as previously – but we show how the Lagrangian construction offers a new way of interpreting it.

Instead of considering $J[u(m)]$ as a functional to minimize on m, we now view J as a function of u only, and accommodate the constraint $L(m)u = f$ by pairing it with a function q in f-space in order to form the so-called Lagrangian[4]

$$\mathcal{L}[u, m, q] = J[u] - \langle q, L(m)u - f\rangle_f.$$

The function q is called Lagrange multiplier in this context, and is arbitrary for the time being. Notice that $\mathcal{L}[u(m), m, q] = J[u(m)]$ regardless of q when $u = u(m)$, i.e., when the constraint is satisfied. This expression can be differentiated to give the desired quantity, namely

$$\frac{d}{dm} J[u(m)] = \langle \frac{\delta \mathcal{L}}{\delta u}, \frac{\delta u}{\delta m}\rangle + \frac{\delta \mathcal{L}}{\delta m}. \tag{4.14}$$

The partials of $\mathcal{L}$ are computed as

- $\frac{\delta \mathcal{L}}{\delta u} = \frac{\delta J}{\delta u} - L^* q$, since $\langle q, Lu\rangle_f = \langle L^* q, u\rangle_u$,
- $\frac{\delta \mathcal{L}}{\delta m} = -\langle q, \frac{\delta L}{\delta m} u\rangle_f$,
- $\frac{\delta \mathcal{L}}{\delta q} = L(m)u - f$.

[4] We reserve the letter $\mathcal{L}$ for the Lagrangian and L for the state equation.

In convex optimization, the traditional role of the Lagrangian is that putting to zero its partials[5] is convenient way of deriving the optimality conditions that hold at critical points, both for the primal and the dual problems. In particular, if the partials of $\mathcal{L}$ are zero at (u, m, q), then m is a critical point of $J[u(m)]$.

The way we make use of the Lagrangian in this section is different, because we aim to derive the expression of $\frac{d}{dm}J[u(m)]$ *away from critical points.* In particular, we are not going to require $\frac{\delta \mathcal{L}}{\delta m}$ to be zero. Still, we find it advantageous to put $\frac{\delta \mathcal{L}}{\delta u} = 0$ as a convenient choice to help simplify the expression of the gradient of J. Indeed, in that case we recover the imaging condition (4.10) from

$$\frac{d}{dm}J[u(m)] = \frac{\delta \mathcal{L}}{\delta m} = -\langle q, \frac{\delta L}{\delta m} u\rangle_f.$$

It is always possible to achieve $\frac{\delta \mathcal{L}}{\delta u} = 0$, by defining q to be the solution of $L^* q = \frac{\delta J}{\delta u}$, the adjoint-state equation (4.9). Note that putting $\frac{\delta \mathcal{L}}{\delta q} = 0$ recovers the state equation $L(m)u = f$.

4.5 Exercises

1. Starting from an initial guess model m_0, a known source function f, and further assuming that the Born approximation is valid, explain how the inverse problem $d = \mathcal{F}[m]$ can be completely solved by means of F^{-1}, the inverse of the linearized forward operator (provided F is invertible). The intermediate step consisting in inverting F is called the *linearized inverse problem.*

   Solution. *Form the incident field as $u_0 = Gf$. Subtract from observed data to get $d - u_0$. Since the Born approximation is assumed valid, we have $d - u_0 \simeq \varepsilon u_1$. Invert for m_1 by solving the system $u_1 = F m_1$, i.e., $m_1 = F^{-1} u_1$. Then form $m = m_0 + \varepsilon m_1$.*

2. Consider the forward wave equation for u_0 in one spatial dimension with an absorbing boundary condition of the form $(\frac{1}{c(0)}\partial_t - \partial_x)u(0) = 0$ at the left endpoint $x = 0$ of the interval $[0, 1]$. Assume that $c(x)$ is locally uniform and equal to $c(0)$ in a neighborhood of $x = 0$.

[5] Or letting 0 be a subgradient of $\mathcal{L}$ in the non-smooth case.

(a) Argue why this choice of boundary condition accommodates left-going waves, but not right-going waves.

(b) Find the corresponding boundary condition for the adjoint-state equation on the backwards field q.

3. *Snapshot migration.* The treatment of reverse-time migration seen earlier involves data $u(x_r, t)$ for an interval in time t, and at fixed receiver points x_r. Consider instead the snapshot setup, where t is fixed, and there are receivers everywhere in the domain of interest. (So we have full knowledge of the wavefield at some time t.) Repeat the analysis of the imaging operator, adjoint to the forward operator that forms snapshot data from singly scattered waves. In particular, find what the adjoint-state wave equation becomes in this case. [Hint: it involves nonzero final data, and a zero right-hand side.]

4. *Sampling.* Call S the linear operator that maps a function $f(x)$ to the vector of point samples $\{f(x_r)\}_r$. Find a formula for S^*. Note that when the linearized forward model F has S as its last operation, then the imaging operator F^* has S^* as its first operation. The presence of S^* explains why we passed from $d_r(t)$ to $d_{\text{ext}}(x, t)$ in the first step of the derivation of the imaging operator.

5. Repeat the general adjoint-state theory by assuming a possibly nonlinear state equation of the form $L(m, u) = f$.