The task then is to choose values for the control variables q, so as to guarantee that the quality feature y is within the performance interval $[y_1,\, y_2]$, for the greatest possible uncertainty in both the linear-model coefficients m and the unknown quadratic coefficients C.

A model of the info-gaps facing the designer is the following extension of eq.(3.56):

$$\mathcal{U}(\alpha, \widetilde{m}, 0) = \{m,\, C: \; |m_i - \widetilde{m}_i| \le \alpha s_i, \; |C_{ij}| \le \alpha, \;\; \forall\, i,\, j\}, \quad \alpha \ge 0 \quad (3.65)$$

The performance requirement is still the interval in eq.(3.57). The robustness is zero if the nominal quality, $\widetilde{m}^T q$, lies outside this interval. If $\widetilde{m}^T q$ is acceptable, then the robustness is evaluated from eqs.(3.59)–(3.61) with the info-gap model of eq.(3.65). Now, unlike eqs.(3.62) and (3.63), we have:

$$\widehat{\alpha}_1(q, y_1) = \frac{\widetilde{m}^T q - y_1}{s^T q + Q} \quad (3.66)$$

$$\widehat{\alpha}_2(q, y_2) = \frac{y_2 - \widetilde{m}^T q}{s^T q + Q} \quad (3.67)$$

where $Q = \sum_{i,j=1}^{3} q_i q_j$. Since Q is positive (since all the control variables q_i are positive) we see that the additional functional uncertainty reduces the robustness. In addition, the robust-satisficing design (which maximizes the robustness) will differ from the case with linear-model uncertainty alone, since now the term $s^T q + Q$ must be reduced, rather than only $s^T q$.

3.2.5 Sequential Decisions

In some decision problems the decision maker is able to measure various features or symptoms and defer decision again and again until the measurements are conclusive. The legendary blind man trying to identify an elephant would do well to defer his decision until he has sampled several distinct features. The central question after each measurement is: how confident will we be if we decide now? Or, conversely, how much could be learned by measuring again? The statistical technique called 'sequential analysis' is the prototype of this type of decision problem [171]. We will use the robustness function to illustrate an info-gap approach, in the context of a very simple example, based on [16].

Let us suppose that we must distinguish between two situations: "normal" or "faulty", "healthy" or "ill", "go" or "no-go", etc. We will sample from a large number, N, of features in order to decide between them. However, the two situations are not uniquely determined by the feature values. Rather, the range of possible values for the feature vector is specified by a different info-gap model for each situation. To keep the example simple let us suppose that these info-gap models are families of N-spheres:

$$\mathcal{U}(\alpha, \widetilde{u}_0) = \left\{u: \; u^T u \le \alpha^2\right\}, \quad \alpha \ge 0 \quad (3.68)$$

$$\mathcal{U}(\alpha, \widetilde{u}_1) = \left\{u : (u - \mathbf{1}_N)^T (u - \mathbf{1}_N) \le \alpha^2\right\}, \quad \alpha \ge 0 \qquad (3.69)$$

where $\mathbf{1}_N$ is the column N-vector whose elements are all 1s. $\widetilde{u}_i$ is the nominal feature vector for situation i, so $\widetilde{u}_0 = 0$ while $\widetilde{u}_1 = \mathbf{1}_N$. Since the uncertainty parameter α is unbounded, any measured collection of features could arise from either situation.

We will sample the first feature, then the second, the third, and so on, until we finally are ready and willing to decide. Let $q(n)$ denote the vector of the n sampled features, for $n = 1, \ldots, N$. $q_j(n)$ is the measured value of the jth feature. $q(n)$ is the decision vector for this problem since the decision to terminate the information-gathering process, as well as the choice of which situation prevails, both depend upon $q(n)$. The question we have posed is: after the nth sample, how reliable is a decision based on the current information? Or, the converse question is: how much additional possible learning are we forgoing by ending the sample and making a decision? If a decision at this point is very reliable then, presumably, there is not too much left to be learned. We will return to a discussion of learning in chapter 8.

After making n measurements we have the sample vector $q(n)$. If we decide to choose between the two situations based on this sample vector, we would do so by comparing it against the two nominal feature vectors $\widetilde{u}_0$ and $\widetilde{u}_1$. Let $\|q(n) - \widetilde{u}_i\|$ denote the Euclidean distance between $q(n)$ and the sub-vector of $\widetilde{u}_i$ containing the corresponding nominal feature values:

$$\|q(n) - \widetilde{u}_i\| = \sqrt{\sum_{j=1}^{n} \left[q_j(n) - \widetilde{u}_{i,j}\right]^2} \qquad (3.70)$$

A decision, when it is made, employs the following 'nearest-neighbor' rule:

$$\text{Decide "0" if and only if } \|q(n) - \widetilde{u}_0\| < \|q(n) - \widetilde{u}_1\| \qquad (3.71)$$

That is, we choose the situation whose nominal feature vector is closest to the measured features.

We now define the robustness function in the spirit of eq.(3.1). $\widehat{\alpha}[q(n)]$ is the greatest value of the uncertainty parameter α which is consistent with not making an erroneous decision based on employing the sample vector $q(n)$ in the algorithm of eq.(3.71). The ambient uncertainty must be sufficiently great to enable the sample vector $q(n)$ to arise from at least one of the info-gap models. Also, as illustrated schematically in fig. 3.7, an error can occur at uncertainty α only if $q(n)$ could come from both models. In the left-hand frame of fig. 3.7 we see that the decision based on the nearest-neighbor rule and sample vector $q(n)$ is "0". At uncertainty α_1 this decision is in fact correct since $q(n)$ cannot arise from $\mathcal{U}(\alpha_1, \widetilde{u}_1)$. However, if the ambient uncertainty is α_2 as on the right-hand side of fig. 3.7, then the same sample vector could have arisen from situation "1", while the nearest-neighbor decision remains the same as before. That is, if the ambient uncertainty is

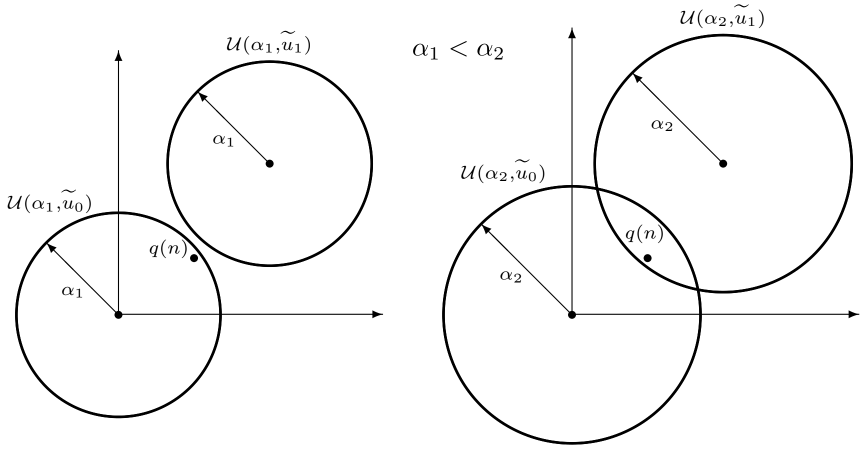

Figure 3.7: Correct decision at uncertainty α_1 (left); possible error at α_2 (right).

large enough so that $q(n)$ belongs to both $\mathcal{U}(\alpha, \widetilde{u}_0)$ and $\mathcal{U}(\alpha, \widetilde{u}_1)$, then the actual situation can differ from the decision according to algorithm (3.71).

From these considerations it is evident that $\widehat{\alpha}[q(n)]$ is the greatest value of α such that $q(n)$ does not belong to both $\mathcal{U}(\alpha, \widetilde{u}_0)$ and $\mathcal{U}(\alpha, \widetilde{u}_1)$:[5]

$$\widehat{\alpha}[q(n)] = \max\{\alpha : \ q(n) \notin \mathcal{U}(\alpha, \widetilde{u}_0) \ \cap \ \mathcal{U}(\alpha, \widetilde{u}_1)\} \tag{3.72}$$

$\mathcal{U}(\alpha, \widetilde{u}_i)$ is a sphere of radius α, so the sampled feature vector $q(n)$ can come from $\mathcal{U}(\alpha, \widetilde{u}_i)$ if and only if:

$$\|q(n) - \widetilde{u}_i\| \le \alpha \tag{3.73}$$

Consequently, for the spherical info-gap models in eqs.(3.68) and (3.69), the robustness — the greatest error-free uncertainty — is simply the greater of the two distances between $q(n)$ and the nominal feature vectors:

$$\widehat{\alpha}[q(n)] = \max\{\|q(n) - \widetilde{u}_0\|, \ \|q(n) - \widetilde{u}_1\|\} \tag{3.74}$$

We now are ready to address the question which we posed earlier: how confident can the decision maker be if the current sample vector is used to make a decision? The question is qualitative so the best we can hope for is a plausible qualitative answer; definitive quantitative answers to qualitative questions simply don't exist. We will briefly consider plausible inference by analogy, to which we will return in detail in chapter 4.

The level α of ambient uncertainty must be great enough for the observed sample vector $q(n)$ to have arisen from at least one of the info-gap models. Consequently, for the spherical models of eqs.(3.68) and (3.69), the lowest possible value of α is:

$$\alpha_{\min} = \min\{\|q(n) - \widetilde{u}_0\|, \ \|q(n) - \widetilde{u}_1\|\} \tag{3.75}$$

[5]We are making a slight notational transgression in eq.(3.72): $q(n)$ is an n-vector while the info-gap models are sets of N-vectors. The meaning of "$q \notin \mathcal{U}(\alpha, \widetilde{u}_0) \cap \mathcal{U}(\alpha, \widetilde{u}_1)$" in eq.(3.72) is that $q(n)$ is not the n-sample vector of any element of the intersection.

If the robustness $\widehat{\alpha}[q(n)]$ — the greatest errorless uncertainty — is 'much greater' than this minimum uncertainty $\alpha_{\min}$, then one can be fairly confident in a decision based on $q(n)$. If $\widehat{\alpha}[q(n)] \gg \alpha_{\min}$ then only an extraordinarily large deviation of the feature vector from its nominal value could lead to an erroneous decision. So the question becomes: what is a 'significant' increment in α? How much greater than $\alpha_{\min}$ must $\widehat{\alpha}[q(n)]$ be, in order to be 'much greater'?

In chapter 4 we will consider several approaches to this problem. As a preliminary treatment of the matter let us consider the following slight generalization of the info-gap models of eqs.(3.68) and (3.69):

$$\mathcal{U}^{\star}(\alpha, \widetilde{u}_i) = \left\{ u : \ (u - \widetilde{u}_i)^T \, (u - \widetilde{u}_i) \le w\alpha^2 \right\}, \quad \alpha \ge 0 \tag{3.76}$$

for $i = 0$ and 1, and where the center-points $\widetilde{u}_i$ are the same as before and w is a positive constant. If $w < 1$ then $\mathcal{U}^{\star}(\alpha, \widetilde{u}_i)$ is a subset of $\mathcal{U}(\alpha, \widetilde{u}_i)$:

$$\mathcal{U}^{\star}(\alpha, \widetilde{u}_i) \ \subset \ \mathcal{U}(\alpha, \widetilde{u}_i), \quad \alpha > 0 \tag{3.77}$$

This means that $\mathcal{U}^{\star}(\alpha, \widetilde{u}_i)$ constrains the uncertain feature vectors u more tightly than $\mathcal{U}(\alpha, \widetilde{u}_i)$ for every positive value of α. In this sense $\mathcal{U}^{\star}(\alpha, \widetilde{u}_i)$ is more informative than $\mathcal{U}(\alpha, \widetilde{u}_i)$. We will discuss informativeness of info-gap models repeatedly in this book.

We will use eq.(3.72) to define the robustness function $\widehat{\alpha}^{\star}[q(n)]$ for the info-gap models in (3.76) just as for the models of eqs.(3.68) and (3.69). These robustnesses are related as:

$$\widehat{\alpha}^{\star}[q(n)] = \frac{1}{\sqrt{w}} \widehat{\alpha}[q(n)] \tag{3.78}$$

That is, enhancing the information (by using info-gap models $\mathcal{U}^{\star}(\alpha, \widetilde{u}_i)$ with $w < 1$) enhances the robustness. Conversely, the increment:

$$\Delta\widehat{\alpha} = \widehat{\alpha}^{\star}[q(n)] - \widehat{\alpha}[q(n)] \tag{3.79}$$

is the amount of robustness which is required in order to compensate for the decrement of information entailed in using $\mathcal{U}(\alpha, \widetilde{u}_i)$ rather than $\mathcal{U}^{\star}(\alpha, \widetilde{u}_i)$.

We will now use $\Delta\widehat{\alpha}$ to address the question of what is a significant increment of robustness.

We wish to know if the robustness, $\widehat{\alpha}[q(n)]$, of a given sample vector $q(n)$ provides strong warrant for terminating the sample. In order to qualitatively evaluate the robustness function we must anchor our judgment in some entity which we understand qualitatively. Let us suppose that we are able to associate qualitative linguistic values to numerical values of w, by indicating for instance that $w = 0.7$ generates a 'substantially' tighter and more informative info-gap model than $w = 1$. In this case, the corresponding robustness increment, eq.(3.79), is 'substantial' or 'significant' *by analogy to* the substantial change in information. An analogy is an inference whereby it is concluded that things which are similar in some attributes will tend to be similar in other attributes as well. The conclusion is plausible, though not definitive, if the two sets of attributes are related, as for instance informativeness of an info-gap model of uncertainty and reliability of a decision employing that model.

We now return to the comparison of the robustness $\widehat{\alpha}[q(n)]$ and the minimum uncertainty $\alpha_{\min}$. Let $\Delta\widehat{\alpha}$ in eq.(3.79) be a 'significant' increment of robustness derived from a significant increment in information. We conclude that $\widehat{\alpha}[q(n)]$ is 'substantially' greater than $\alpha_{\min}$ and, therefore, a decision based on $q(n)$ is dependable, if:

$$\widehat{\alpha}[q(n)] > \alpha_{\min} + \Delta\widehat{\alpha} \tag{3.80}$$

If this relation does not hold, then the sample should be extended by measuring more features.

In chapter 4 we will pursue a more thorough study of reasoning by analogy and qualitative calibration of increments of robustness.

3.2.6 Project Scheduling with Uncertain Task Durations

The management of a complex multi-task project requires decision-making in many different areas. One of the more quantitative managerial decisions, for which info-gap decision-support tools can be useful, is the design of the flow chart of task execution. A simple flow chart is shown in fig. 3.8. Each box represents a task and the arrows indicate the order of implementation. While all the tasks must be completed, and while many constraints exist on the task-sequencing, alternative flow charts are usually possible for the same project, sometimes based on different technologies for task implementation and incurring different costs. The problem facing the manager is to choose the best among the possible alternatives. The difficulty in making such a choice is that the factors which may retard the implementation of the tasks are often highly uncertain. Some tasks may involve engineering R & D whose resource requirements are hard to estimate. Some tasks may be vulnerable to unknown external factors — economic anomalies, natural catastrophes, political developments — which are impossible even to name or enumerate. In order to choose from among alternative flow charts, the manager needs a concise assessment of the project reliability against highly uncertain adversity. The robustness function provides precisely this.

We will consider a simple example in this section, and return to study various other aspects of the problem later. In section 5.4 we will study both robustness and opportuneness functions in attempting to choose between alternative task flow charts. In section 7.5 we will examine the enhanced robustness of an entire project which results from gathering information about a highly uncertain task.