Normalization in Multi-Objective Optimization

Helmut Mausser, Algorithmics Inc.

Consider the multiple-objective constrained optimization problem

$$\begin{array}{ll} \min & \sum_i w_i f_i(\boldsymbol{x}) \\ \text{s.t.} & \boldsymbol{x} \in \Omega \end{array}$$

where

- $\Omega \subseteq R^n$ is convex
- f_i: $R^n \rightarrow R$ is linear or quadratic
- $w_i \in R$ is the weight of the ***i***-th objective

In the context of financial optimization, $\boldsymbol{x}$ represents a vector of position sizes for a set of n securities and Ω is typically defined by a set of linear constraints. The objectives may include any of a number of portfolio characteristics, such as

- variance of return
- tracking error (variance of excess return relative to a benchmark)
- expected return
- trading cost
- penalties associated with various soft constraints relating to the portfolio's composition

Generally, objectives compete, i.e., improvement in one objective is obtained only at the expense of others. A key input to a multi-objective problem is the relative importance, or priority, attached to an objective by the decision-maker (DM). A priority is a real number u, where importance increases with the magnitude of u (the DM might specify priorities as "low/medium/high," with each subjective level subsequently being translated to a predefined numerical value).

The concept of importance must be defined in a manner that is relevant to the DM. One intuitively appealing approach associates importance with the relative deviation from some "ideal" value (sometimes called an "aspiration level" or a "reference point"). Thus, in an optimal solution, competing objectives that are deemed equally important should attain approximately the same level, in percentage terms, of their respective ideals, while a higher priority objective should be closer to its ideal than a competing lower priority objective. Conversely, if objectives do not compete then all may attain their ideals regardless of their respective priorities.

For optimization purposes, priorities are converted to weights, which must account for differences in the magnitudes of the objectives. For example, if objectives i and j represent the variance of return and the trading cost, respectively, then $f_i(\boldsymbol{x})$ may be of order 10^{-2} while $f_j(\boldsymbol{x})$ may be of order 10^4. Suppose the objectives are considered to be equally important, and that each is assigned a priority of u by a DM. In this case, it is necessary that $w_i = 10^6 w_j$ and so the priorities provided by the DM must be "normalized" in a consistent manner, e.g., $w_i = u$, $w_j = 10^{-6} u$.

Various normalization schemes have been proposed, most of which rely on advance knowledge of the ideals. For example, if f_i^* is the ideal value of $f_i(\boldsymbol{x})$ then the weight w_i can be obtained from the priority u_i as follows:

$$w_i = \frac{u_i}{f_i^*}$$

This approach is practical when the ideal is a "goal" specified by the DM (e.g., $f_i(\boldsymbol{x})$ is a soft constraint which penalizes deviations from some constant target f_i^*). The difficulty occurs when f_i^* is not provided by the DM. Suppose, for example, that $f_i(\boldsymbol{x})$ is an objective term that equals the portfolio's tracking error (which is to be minimized). In this case, f_i^* can be obtained by solving a separate single-objective problem. However, the computational effort required to solve one or more such problems may make this approach impractical from a performance standpoint. In particular, Algorithmics' optimization functionality is used on a "real time" basis and thus a turnaround time of no more than several minutes is required.

Thus, the key question is whether or not one can quickly and reasonably approximate f_i^* using only information that is available when the problem is formulated. Such information includes.

- coefficient data
- constant targets for soft constraints
- trading limits (i.e., bounds on the number of units that may be bought or sold for each tradable security)
- initial holdings, if any, of tradable securities

Note that the objective is not to construct an entire efficient frontier of all non-dominated solutions; we simply want to find a solution that adequately reflects the importance weights assigned by the DM. While iterative procedures exist for traversing the efficient frontier to obtain the most desirable solution for the DM, the practicality of such an approach again depends critically on the computational time.

One technique, which has shown promise in testing to date, is to approximate f_i^* by $f_i(\boldsymbol{x^0})$, where $\boldsymbol{x^0}$ denotes the initial holdings in the portfolio. However, this fails if $f_i(\boldsymbol{x^0}) = 0$, as is typically the case when there are no initial holdings.

Reference

Korhonen, P. (1998), "Multiple Objective Programming Support," IIASA Report IR-98-010.