We developed an algorithm to perform patient localization in three dimensions (3D) using *a priori* 3D information of the patient (a prescription CT dataset) and a series of N two-dimensional (2D) projection radiographs (where $N \geq 1$) obtained at treatment time. Patient and site-specific anatomic features in the prescription CT dataset are selected and identified in the projection radiographs. The 3D coordinates of these features, their corresponding 2D projected coordinates, along with the motion characteristics of the non-rigid imaging system, are used in an optimization algorithm to determine the proper rigid-body transformation (translation and rotation) between the actual patient position at treatment time and the patient position in the prescription CT dataset. This transformation can then be used to apply a patient setup error correction and/or adjust treatment delivery parameters. Results obtained from ball bearing and humanoid phantoms show that localization can be determined with a precision of 0.4 mm (along each axis) for the translations and 0.1° (about each axis) for the rotations when the non-rigid motion characteristics of the imaging system are taken into account. When the imaging system is assumed to be rigid and ideal, localization precision is 1.3 mm for the translations and 0.2° for the rotations. However, localization inaccuracy (where the non-rigid imaging system is the truth) is as large as 7.1 mm for the translations and 0.4° for the rotations, when using open-field projection radiographs (free of treatment port information).