FAQs for the Practical Use of MIM SureCalc® MonteCarlo

Document History

Version	Date	Description	Editor
1	14.02.2019	Initial Version	MLA
2	20.02.2019	Amendments	MLA
3	22.02.2019	Adding figure about arc-discretization effects	MS
4	27.02.2019	Added section on Gamma, added suggestions by proof-readers (physics)	MLA

What is the impact of voxel size on accuracy?

Particle transport through a voxelized geometry has two effects: the material is constant throughout the voxel and the dose value is the average of all energy depositions inside the voxel.

The input computed tomography (CT) image is usually downsampled to the calculation voxel size, which leads to a certain blurring of density interfaces, for example around the patient outline, internal airways, lung tissue and bones. The resampling preserves mass, i.e. it has no effect on the total absorption of energy, but may have an effect for voxels that contain very different tissues. The natural limit for this effect is the original CT resolution (both determined by image pixel size, slice spacing and reconstruction kernels). Using smaller voxels than image resolution for dose computation will have no beneficial effect on accuracy with respect to density effects.

Although the particle transport can deposit energy in any position within one voxel, the dose value for a voxel is an average over the voxel volume. This results in a blurring of the dose distribution. In contrast to the density blurring, there is no lower limit for the voxel size where this effect disappears. However, due to the physics of photon and electron transport, voxel sizes below 1 mm do not improve the spatial resolution of the dose distribution further. Observe that a similar dose averaging effect is also produced by the finite volume of dose detectors. For example, for water phantom measurements with a 0.125 ccm ionization chamber, a voxel size of 5x5x5 mm^3 would produce a similar dose blurring effect.

In short, there is no rationale for voxel sizes below 1 mm. For all clinical applications apart from extremely small stereotactic volumes, a voxel size of 2 mm may be considered as more than adequate.

What is the meaning of the statistical uncertainty?

Monte Carlo methods estimate the dose in a voxel up to some residual uncertainty. SciMoCa employs uncertainty levels of "Extra Fast", "Fast", "Precise", "Extra Precise", which correspond to uncertainties of 4%, 2%, 1%, and 0.5% respectively. These uncertainty values are to be understood as random statistical uncertainties ("noise") and describe the average uncertainty of all voxels with D > 0.7 * Dmax. Notice that the dose uncertainty increases with 1/D, i.e. it is higher in low dose regions than in the target. Hence, the uncertainty settings are for the minimum relative uncertainty. The noise of the dose distribution is approximately Gaussian; this means that at "Extra Precise", only 70% of the voxels with dose greater than 70% of Dmax will have an uncertainty of less than 0.5%. However, among 10000 voxels of a target volume, 455 will have a deviation of more than 1% and 27 of more than 1.5%.

MIM Software Inc. • 25800 Science Park Drive - Suite 180, Cleveland, Ohio 44122 • 866-421-2536

Monte Carlo uncertainties have an impact on DVHs as well. This is most visible for target volumes and affects the maximum and minimum dose most. Target volume mean dose, as well as normal tissue dose-volume points, will be much more stable.

SciMoCa´s uncertainty levels "Extra Fast", "Fast", "Precise", "Extra Precise" mean that the number of histories from step to step is quadrupled. Typically, simulations require between 10^8 and 10^{10} histories, depending on voxel size, target volume and uncertainty setting.

By default, repeated runs of the same case would yield the same results within floating point precision, because the seeds of the random number generators (RNGs) are fixed (each computing thread has an individual seed). It is also possible to set an option which starts the RNGs with random seeds at each run, in which case the results would differ within statistical uncertainty. Notice further that the results depend on hardware concurrency, i.e. the number of threads that can be started on a computer.

What is the meaning of Dose-To-Water versus Dose-To-Medium calibration?

Dose-To-Medium (D2M) is the natural result of any Monte Carlo computation. By virtue of the application of cavity theory, it is possible to convert it to Dose-To-Water (D2W). A proper re-calibration requires the knowledge of the secondary electron spectrum in each voxel, which is not normally stored in clinical Monte Carlo computations for efficiency reasons. Thus, the customary way (and the one employed in SciMoCa) is the application of a material specific, voxel-wise calibration factor, which is an approximation. Because of this, and because the D2W calibration amplifies differences in material definition, dose computation algorithms should always be compared on the basis of D2M calibration. This is particularly true for phantom computations, where material compositions may differ even more between algorithms than in tissues.

When comparing Monte Carlo results to measurements, the situation is also not clear cut. Even though a measurement device may be calibrated to D2W, and may be considered "water-equivalent" while in water, the situation changes when it is immersed in a different material. Normally, the device is neither water-equivalent with respect to its stopping power, nor is it small enough to be considered without influence on the secondary electron flux. Therefore, detectors do not measure dose-to-water when immersed in non-water material. If a detector has a spectral-, depth-, off-axis- or beam quality dependence, it will also show a deviation from dose-to-water when measuring in non-water material. Thus, differences between simulation and measurements will be smallest in mid-sized, flattened beams and largest in stereotactic flattening-filter free beams.

Which materials are identified?

SciMoCa assigns material properties (cross sections, stopping powers) according to the mass density values derived from the CT image. The materials are air (<0.01 g/cm^3), lung (<0.75 g/cm^3), fat-like soft tissue (<0.99 g/cm^3), soft tissue (<1.12 g/cm^3), water (=1.0 g/cm^3), bone (>1.12 g/cm3), titanium (=4.54 g/cm^3)and steel (=8.0 g/cm^3). Material properties within these categories are modified according to mass density, i.e. soft tissue with density 1.01 and 1.05 are of the same type with respect to atomic composition, but have different material properties.

How does SciMoCa compare to other Monte Carlo codes?

General purpose Monte Carlo codes, like EGSnrc, MCNP or PENELOPE, can be applied to a much greater range of problems than specialized codes like SciMoCa. Although they share a great many methods, SciMoCa omits complexity that does not pay off for the case of megavoltage, external beam irradiation of patients. For example, SciMoCa is not designed for kilovoltage radiation or ionization chamber simulations. The three main aspects of specialization are:

1. the restriction to rectilinear voxel geometries, enabling faster tracking of particles
2. the restriction to few material types, with individual material composition, enabling faster material property look-up and on-the-fly computation
3. the omission of interactions at low photon and electron energies, like spin effects or atomic relaxation, enabling a simpler handling of material properties

SciMoCa shares these restrictions with other specialized Monte Carlo codes that have been introduced clinically, like VMC++ and XVMC. Like these, it derives its fundamental concepts from the works of Bielajew, Fippel, Kawrakow and Rogers and is thus in the tradition of the EGSnrc/XVMC/VMC++ family of codes as a class II condensed history algorithm with sophisticated variance reduction techniques (VRTs).

The SciMoCa patient transport code is a thorough implementation of these physics concepts, with some additional accelerations afforded by exploiting contemporary CPU architectures. It was constructed against the benchmark of EGSnrc with the goal of keeping deviations in the toughest artificial situations to a maximum of 2%, in order to achieve agreement within statistical uncertainty in clinical situations (see figures 1 and 2)

MIM Software Inc. • 25800 Science Park Drive - Suite 180, Cleveland, Ohio 44122 • 866-421-2536

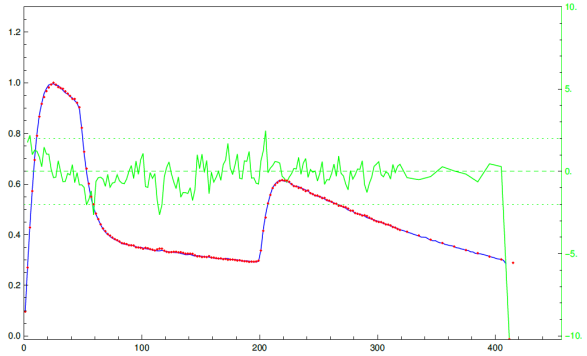

Fig 1: 10x10 mm² field, 6 MeV mono-energetic point source, 150 mm slab of ICRU-lung material with 0.25 g/cm³ starting at 50 mm depth. Blue: SciMoCa, Red: EGSnrc. Residual errors at interfaces are a consequence of simplified boundary crossing algorithm in SciMoCa.

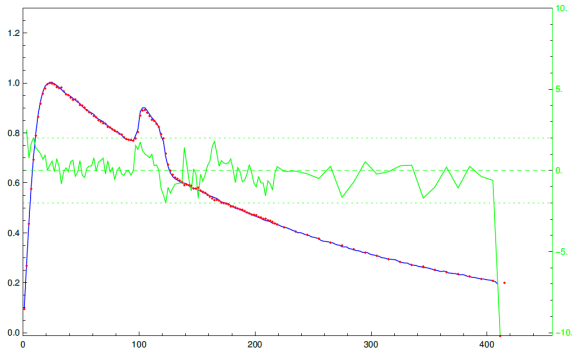

Fig 2: 10x10 mm² field, 6 MeV mono-energetic point source, 20 mm slab of ICRU-bone material with 2.00 g/cm³ starting at 50 mm depth. Blue: SciMoCa, Red: EGSnrc. Residual errors at interfaces are a consequence of simplified boundary crossing algorithm in SciMoCa.

The main discriminator to the EGSnrc/XVMC/VMC++ family is the virtual source accelerator head and collimator model ("virtual source model", VSM), which makes heavy use of VRTs for speed-up. SciMoCa´s unique feature is the elaborate tuning of the VRTs across source model and patient transport code, in order to maximize efficiency[1] and minimize non-Gaussian noise.

Why does SciMoCa not use a GPU architecture?

For the efficient simulation of the accelerator head and variable collimators, variance reduction techniques (VRT) can be applied to great benefit. In principle, VRTs are ways to perform computations with fewer operations, very much like a*u + a*v = a*(u + v) saves one operation. Usually, VRTs come with some overhead in terms of memory management, and with some price in terms of code complexity. Both factors are a greater problem for a GPU architecture, which is specialized for executing the same set of instructions in parallel on many data entities (SIMD, single instruction, multiple data) than for a CPU architecture (MIMD, multiple instructions, multiple data). SciMoCa is designed to be versatile and work with very diverse accelerator head geometries, and has a Precisely balanced suite of VRTs to couple the accelerator head model with the patient model. Hence, the code executes preferentially on a CPU architecture.

Why does SciMoCa use Virtual Source Models?

Virtual sources are an analytical representation of a phase space. Compared to phase space files with a finite number of particles, they have the advantage of high computation speed (because of minimum memory throughput) and zero latent variance (because particles do not have to be re-used to arrive at the required number of histories). They also offer an unrivalled degree of flexibility, which is essential for wide-spread clinical deployment for a large number of different linacs.

As in all aspects of Monte Carlo simulations, it is important to hit the right balance between simplification and accuracy. SciMoCa´s virtual sources are constructed from phase space files generated by BEAMnrc. A unique process extracts the numbers and properties of up to five sources from the phase spaces. Sources can be planar or volumetric, have angular-dependent spectra and energy fluence and can contribute photons, electrons or back-scatter into the monitor chamber. For each accelerator type, there is a characteristic set of source properties and parameters that are fixed, and some parameters that need to be fitted. In result, the virtual sources have the potential to reproduce any phase space of an idealized linac in terms of its dosimetric accuracy, but have the added flexibility to be easily tuneable to a real linac.

The beam modelling process would generate a number of reference beam models that cover the variability of beam tunes in practice. For example, for a given linac type, the maximum energy of a 6 MV beam quality may vary between 5.8 MeV and 7.2 MeV. This family of reference models is the basis for a similarity-search-and-interpolation strategy that produces a specific beam model as a new family member.

How does SciMoCa simulate variable collimators?

SciMoCa employs a unique generic 3D collimator model which can account for such details as single- or double focussing leaf design, stepped leaf sides, rounded leaf and jaw faces, inter-leaf leakage etc. The collimator elements are not described to the last detail (currently, 16 different MLC types are supported), but well enough to be dosimetrically indistinguishable from full simulations. Because of inaccuracies in vendor drawings, manufacturing tolerances and unknown or variable material compositions, the models may need to be tuned to individual linacs, even though they should be invariable. Scattered photons are suppressed if they would not contribute to the dose distribution. Secondary electrons are suppressed.

MIM Software Inc. • 25800 Science Park Drive - Suite 180, Cleveland, Ohio 44122 • 866-421-2536

How does SciMoCa simulate dynamic deliveries?

In contrast to analytical algorithms (Pencil Beam, Convolution/Superposition, Collapse Cone, Linear Boltzmann Equation Solvers), Monte Carlo does not have to discretize the time dimension. Instead, the time axis can be sampled just like any other dimension of a simulation: for example particle energy, direction or origin. In practice, every particle is assigned a random time of emission. During the simulation, the collimators, gantry and couch are placed according to this emission time. Hence, every particle history samples a different time point and is tracked through a different linac configuration.

A typical example is the discretization of the gantry angle of VMAT-arcs as done as approximation by analytical algorithms, which manifests in characteristic dose difference patterns w.r.t. to the dose calculated as 'continuous arc' by Monte Carlo, see fig. 3.

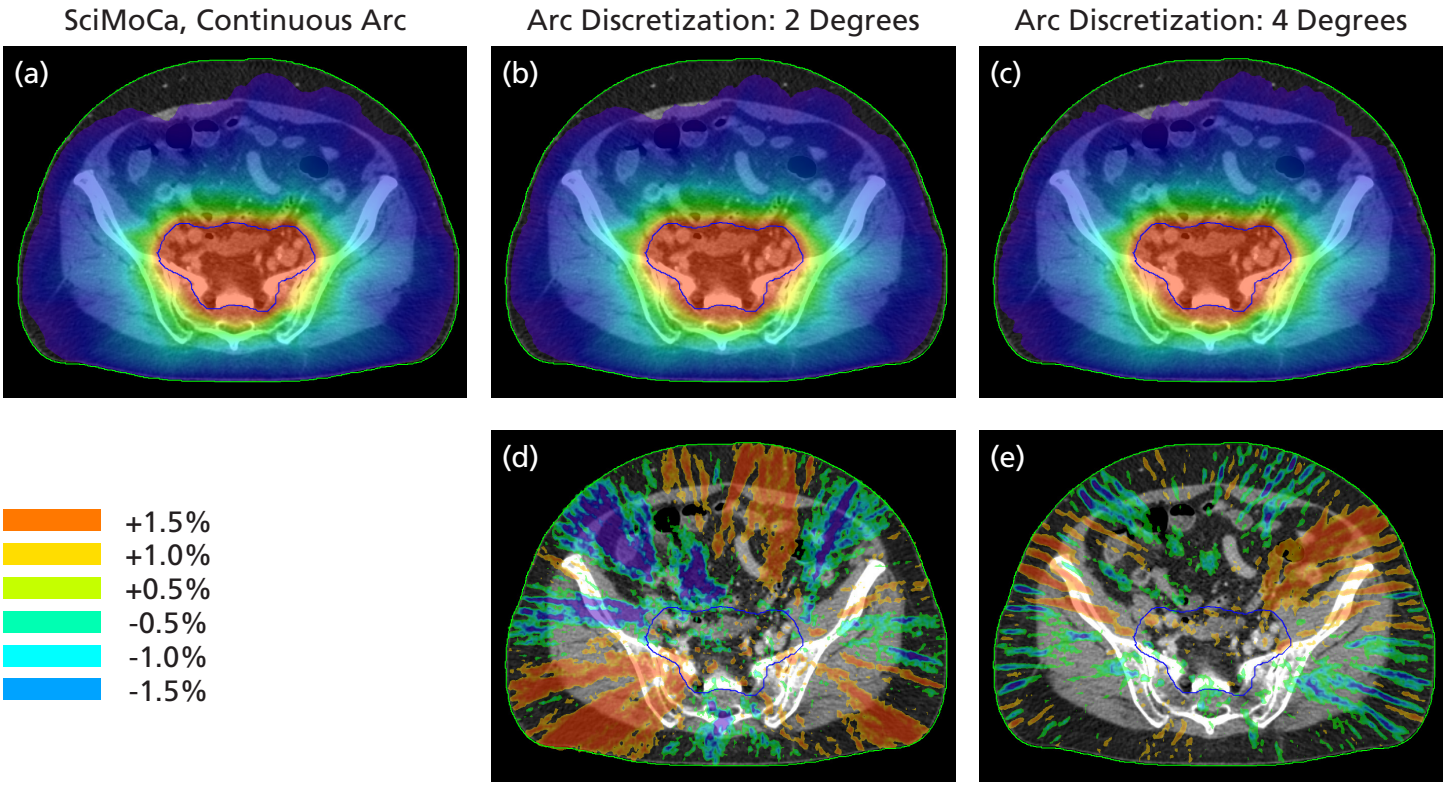

Fig 3: (a) SciMoCa dose (with continuous sampling of the gantry angle) calculated for a 15MV-rectum VMAT plan; (b) and (c) show effects of arc discretization into 2- or 4-degree sectors, respectively.

(d) and (e) show the relative dose differences to the SciMoCa-dose in (a).

MIM Software Inc. • 25800 Science Park Drive - Suite 180, Cleveland, Ohio 44122 • 866-421-2536

How was SciMoCa validated?

Throughout its development, SciMoCa patient transport code was benchmarked against EGSnrc with tests designed to highlight any weakness in the algorithm. These tests largely exceeded the demands of clinical cases, but had simplified beam properties. Typically, they would be point sources, rectangular fields and mono-energetic beams. SciMoCa was validated for materials between mass density 0.25 g/cm^3 and 8.0 g/cm^3 and photon energies between 1 MeV and 25 MeV. The maximum deviation for all test cases was 2%, but was typically not detectable within statistical uncertainty. For this reason, the validation of the entire algorithm can concentrate on the simulation of the radiation source and collimators.

For the validation of the linac beam models, a strong preference was given to 1D measurements under controlled conditions, preferentially in a water phantom. Measurements with film, 2D/3D detector arrays and in solid phantoms can quickly result in uncertainties of 2%, which are not sufficient for algorithm validation. Measurements of realistic plans served the purpose of comprehensive end-to-end tests after the validation of the beam models were performed.

Beam model validation can only be performed for a specific linac, however the variability between linacs of the same type is typically greater than the accuracy limits of SciMoCa´s validation tests. In order to achieve a consistently high performance of SciMoCa, individual beam modelling and validation is indispensable.

How should SciMoCa be validated by customers?

Each beam model is validated with the water phantom data provided for its commissioning. A report summarizes the agreement. There may be some measurements that deviate more than others, and the former would be a good starting point for validation against additional water phantom measurements. Typically, situations that are challenging for beam modelling are also challenging for measurements, so that the presence of systematic measurement errors may quickly yield a very confusing picture.

One fundamental advantage of Monte Carlo over any other dose computation algorithm is its modularity - the source properties do not change with collimator settings - which results in self-consistency of the simulated doses. For example, the same photons that hit the collimators for a 100x100 mm^2 field will hit the collimators for a 50x50 mm^2 field. Self-consistency in this context means: it is not possible that the larger field is correct and the smaller field is wrong in the simulations, and vice versa. This statement typically holds for field sizes that differ less than a factor 2. With this rule of thumb, measurement data can be identified as problematic and inconsistent with the rest. It is also possible to modify the initial experimental plan such that the feature in question can still be investigated, but with smaller experimental error. For example, the depth dose curve of a 10x10 mm^2 and a 30x30 mm^2 field contain the same information with respect to the beam model (the primary spectrum), but the latter is substantially easier to measure to an accuracy that allows validation of SciMoCa.

Each SciMoCa beam model is uniquely tuned to one individual linac, but it is also derived from a family of realistic reference beam models, which lends additional robustness to the performance. Therefore, in practice, the dominating source of disagreement between SciMoCa, TPS and measurements are small field output factors and MLC calibration. A solid calibration of the MLC positions in the Monte Carlo simulation is virtually the most important prerequisite for accuracy of dynamic IMRT/VMAT treatment plans. MLC calibration can be determined from static abutting field strips or via a dosimetric leaf gap (DLG) experiment. Notice also, that due to manufacturing tolerances, the inter-leaf leakage of some MLC models can vary substantially, so that this parameter may have to be tuned.

Typically, the deviations caused by MLC calibration / small field output factor errors become worse with a higher modulation degree of the treatment plans. Although different errors can compensate each other, there should be a trend in target dose and/or the most spared normal tissues with increasing number of MU/Gy. Thus, it is advisable to start SciMoCa/TPS/measurement cross-validation with low modulation cases, maybe even the reference field, and gradually move towards more demanding plans.

MIM Software Inc. • 25800 Science Park Drive - Suite 180, Cleveland, Ohio 44122 • 866-421-2536

Further, a very important point is to ensure the correct configuration setup: is the Hounsfield-to-Density calibration correct? Are mass- or electron densities used in the calibration table? Is the handling of the external patient contour correct? Is the handling of the treatment table correct? Is the monitorunit calibration correct?

What are the correct density values for phantom materials?

Phantom materials, especially when they are not labelled as "water equivalent", pose a double problem for accurate dosimetric validation. Assignment of a density according to a Hounsfield calibration table will almost certainly yield the wrong result, as will a measurement without a material- and detector specific calibration factor. Furthermore, SciMoCa will interpret non-water materials as "tissue with mass density x", which causes errors especially for lung- and bone-mimicking phantom materials.

It is therefore of paramount importance to determine both the specific mass density and the detector calibration factor for every dose measurement in a solid phantom. Notice that these parameters can depend on beam quality (energy AND flattening filter presence), field size (IMRT/VMAT vs. 3D conformal vs. SBRT) and measurement depth, particularly for measurement in non-water equivalent phantom materials. Depending on the planned experimental setup, these parameters can deviate from manufacturer recommendations.

A good way to proceed is to determine the appropriate phantom mass density to be used for SciMoCa-simulations by tweaking it until a phantom measurement and its Monte Carlo simulation match perfectly for a mid-sized field (between 50x50 and 100x100 mm^2). Notice, since the detector calibration needs to be determined in parallel, the measurement should include multiple points for depths between the dose maximum and about 100 mm. Modify the mass density until all dose points deviate by a constant factor, which is the detector calibration.

As a rule of thumb, effective mass densities for SciMoCa are lower than real mass densities, because most polymers have a Z_{eff} that is slightly lower than the Z_{eff} of a tissue of the same mass density. For example, for a 6 MV beam, PMMA has an effective mass density of 1.13 (vs. 1.145 real) and a calibration factor for 0.125 ccm ionization chamber of roughly 1.01 (Chamber measures more than SciMoCa D2M). Larger correction factors will be obtained for lung- and bone-mimicking materials.

What is there to know about Gamma analysis for Monte Carlo dose computations?

In principle, the Gamma analysis should be independent of such parameters as voxel size or grid alignment - interpolation of the reference and compare grid should be the standard in any interpolation. However, Gamma analysis can be perturbed by statistical uncertainties in the reference grid and compare grid in a different manner[2]. A noisy reference distribution artificially increases the spread of Gamma-values (thereby diminishes pass rates), while a noisy compare distribution artificially decreases Gamma-values (thereby increases pass rates). Notice that also analytical algorithms produce discretization residuals in the dose distribution (they are not as smooth as they would physically be) which can have effects on Gamma pass rates and can influence the performance of dose interpolation.

MIM Software Inc. • 25800 Science Park Drive - Suite 180, Cleveland, Ohio 44122 • 866-421-2536

Footnotes:

1. Efficiency in Monte Carlo simulations is often expressed as $e = 1/(s^2T)$, where T is the computation time and s the statistical uncertainty. While this is the global view, it is also possible to compare the efficiency per history $_{eh} = 1/(s^2N)$, where N is the number of histories. Variance reduction techniques aim to increase eh, but come at an expense of time per history T/N. Notice that eh is independent of the specific hardware, but depends on the hardware architecture (e(CPU) > e(GPU)). In contrast, e can be boosted with investment in more powerful hardware.
2. An excellent summary of the mathematics of Gamma analysis and the influence of statistical uncertainties and finite grid size can be found in Clasie et al., Phys. Med. Biol. 57, 6981-6997 (2012).

Feature	Value/Reference	Description
electron cut-off energy for last Multiple Scatter step	< 240 keV	
fractional energy loss of electron Multiple Scatter step	0.12	
bremsstrahlung production cut-off energy	> 6 keV	
photon cut-off energy (local energy deposit)	< 60 keV	
minimum/maximum particle weight(Russian Roulette ratio)	0.5 < w < 2.0	
maximum photon energy	< 25 MeV	
KERMA-approximation threshold energy	< 1.0 MeV	
Material properties	ICRU 46	XVMC
Material property computation	Kawrakow 1996, Fippel 1999	VMC, XVMC, VMC++
Photon effects	Photoelectric absorption, Compton scatter, Pair production (Kawrakow 2000a)	XVMC, VMC++
Electron effects	Elastic scatter, Møller, Bremsstrahlung (Kawrakow 1996, 2000a)	XVMC, VMC++
Positron effects	Elastic scatter, Bhabha, Bremsstrahlung(Kawrakow 1996, 2000a)	XVMC, VMC++
Multiple Scatter theory	Kawrakow 2000b	EGSnrc, VMC++
Multiple Scatter boundary crossing	Kawrakow 1997, 2001	XVMC, VMC++
Variance reduction techniques	Woodcock tracking, adaptive history repetition, adaptive particle splitting, Russian Roulette, KERMA-approximation	XVMC, VMC++
Random numbers	Pseudo, cycle 248-1	

Fippel 1999: M. Fippel: Med. Phys. 26, 1466 (1999)

Kawrakow 1996: I. Kawrakow, M. Fippel, K. Friedrich: Med. Phys. 23, 445 (1996)

Kawrakow 1997: I. Kawrakow: Med. Phys. 24, 505 (1997)

Kawrakow 2000a: I. Kawrakow, M. Fippel: Phys. Med. Biol. 45, 2163 (2000)

Kawrakow 2000b: I. Kawrakow: Med. Phys. 27, 485 (2000)

Kawrakow 2001: I. Kawrakow in: A. Kling et al. (eds.), Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications, Springer-Verlag Berlin Heidelberg (2001)

MIM Software Inc. • 25800 Science Park Drive - Suite 180, Cleveland, Ohio 44122 • 866-421-2536