



King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Windows-Yule, C. R. K., Herald, M. T., Nicusan, A. L., Wiggins, C. S., Pratz, G., Manger, S., Odo, A. E., Leadbeater, T., Pellico, J., T. M. de Rosales, R., Renaud, A., Govender, I., Carasik, L. B., Ruggles, A. E., Kokalova-Wheldon, T., Seville, J. P. K., & Parker, D. J. (in press). Recent Advances in Positron Emission Particle Tracking: A Comparative Review. *REPORTS ON PROGRESS IN PHYSICS*.

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

ACCEPTED MANUSCRIPT

Recent Advances in Positron Emission Particle Tracking: A Comparative Review

To cite this article before publication: Christopher R. K. Windows-Yule *et al* 2021 *Rep. Prog. Phys.* in press <https://doi.org/10.1088/1361-6633/ac3c4c>

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is © 2021 IOP Publishing Ltd.

During the embargo period (the 12 month period from the publication of the Version of Record of this article), the Accepted Manuscript is fully protected by copyright and cannot be reused or reposted elsewhere.

As the Version of Record of this article is going to be / has been published on a subscription basis, this Accepted Manuscript is available for reuse under a CC BY-NC-ND 3.0 licence after the 12 month embargo period.

After the embargo period, everyone is permitted to use copy and redistribute this article for non-commercial purposes only, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by-nc-nd/3.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions will likely be required. All third party content is fully copyright protected, unless specifically stated otherwise in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

Review Article

Recent Advances in Positron Emission Particle Tracking: A Comparative Review

C.R.K. Windows-Yule¹, M.T. Herald¹, A.L. Nicuşan¹, C.S. Wiggins^{2,3}, G. Pratz^{4,5}, S. Manger¹, A.E. Odo^{6,7}, T. Leadbeater⁷, J. Pellico⁸, R.T.M. de Rosales⁸, A. Renaud⁹, I. Govender¹⁰, L.B. Carasik², A.E. Ruggles¹³, Tz. Kokalova-Wheldon¹⁴, J.P.K. Seville¹, D.J. Parker¹⁴

¹School of Chemical Engineering, the University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK

²Department of Mechanical and Nuclear Engineering, Virginia Commonwealth University, 401 West Main Street, Box 843015, Richmond, Virginia 23284

³Department of Physics and Astronomy, University of Tennessee, Knoxville, 1408 Circle Drive, Knoxville, TN 37996

⁴Department of Radiation Oncology, Division of Medical Physics, Stanford University School of Medicine, Stanford University, Stanford, CA, USA

⁵Molecular Imaging Program at Stanford (MIPS), School of Medicine, Stanford University, Stanford, CA, USA

⁶Department of Physics, Federal University Oye-Ekiti, Nigeria

⁷Department of Physics, University of Cape Town, Rondebosch, 7701, South Africa

⁸School of Biomedical Engineering & Imaging Sciences, King's College London, St. Thomas' Hospital, London SE1 7EH, UK

⁹School of Mathematics, the University of Edinburgh, Old College, South Bridge, Edinburgh EH8 9YL

¹⁰School of Engineering, University of KwaZulu Natal, Glenwood, 4041, South Africa

¹¹Mintek, P/Bag X3015, Ranburg, Gauteng 2121, South Africa

¹²Centre for Minerals Research, University of Cape Town, P/Bag Rondebosch, 7701, South Africa

¹³Department of Nuclear Engineering, University of Tennessee, Knoxville, 1412 Circle Drive, Knoxville, TN 37996

¹⁴School of Physics and Astronomy, the University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK

E-mail: c.r.windows-yule@bham.ac.uk

May 2021

Abstract. Positron emission particle tracking (PEPT) is a technique which allows the high-resolution, three-dimensional imaging of particulate and multiphase systems, including systems which are large, dense, and/or optically opaque, and thus difficult to study using other methodologies. In this work, we bring together researchers from the world's foremost PEPT facilities not only to give a balanced and detailed overview and review of the technique but, for the first time, provide a rigorous, direct, quantitative assessment of the relative strengths and weaknesses of all contemporary PEPT methodologies. We provide detailed explanations of the methodologies explored,

CONTENTS

including also interactive code examples allowing the reader to actively explore, edit and apply the algorithms discussed. The suite of benchmarking tests performed and described within the document is made available in an open-source repository for future researchers.

Keywords: Positron emission particle tracking, PEPT, positron imaging, nuclear imaging, biomedical imaging, granular materials, particulate media, multiphase systems

Submitted to: *Rep. Prog. Phys.*

Contents

1	Introduction	3
2	The Fundamentals of PEPT	4
2.1	PEPT in an Ideal World	5
2.2	PEPT in the Real World	6
3	PEPT Hardware	7
3.1	PEPT Detectors	7
3.1.1	An Overview of Detector Physics–	7
3.1.2	Conventional Detector Systems (Ring and Planar Scanners)–	8
3.1.3	Modular Cameras–	10
3.1.4	Evolution of PEPT Camera Capabilities–	11
3.2	PEPT Tracers	11
3.2.1	Fundamental Concepts	12
3.2.2	Direct Activation	14
3.2.3	Indirect Activation	15
4	Simulating PEPT	16
4.1	The Value of Simulation	16
4.2	Modelling PEPT Hardware	16
4.3	Simulating a PEPT Experiment	17
5	PEPT Methodologies – an Introduction	18
5.1	The Birmingham Method	19
5.2	The Line-Density Method	21
5.3	G-Means Clustering	22
5.4	Spatiotemporal B-spline reconstruction (SBSR)	25
5.5	The Feature Point Identification Method (FPI)	26
5.6	Voronoi-Based Multiple Particle Tracking (VMPT)	29

<i>CONTENTS</i>		3
5.7	The Spherical Density Method	29
5.8	PEPT using Machine Learning (PEPT-ML)	33
5.9	PEPT using Expectation Maximisation (PEPT-EM)	34
6	PEPT Methodologies – a Quantitative Comparison	37
6.1	A Standard Testing Framework for PEPT	38
6.2	Single Particle Tracking Tests	39
6.2.1	Minimum activity test–	39
6.2.2	Maximum velocity test–	40
6.2.3	Scatter sphere test–	43
6.2.4	Field of view test–	46
6.3	Multiple Particle Tracking Tests	50
6.3.1	Minimum separation test–	50
6.3.2	False positive test–	52
6.3.3	Orientation tracking test–	52
6.3.4	Large tracer number test–	55
7	PEPT Applications	56
7.1	Current Applications	57
7.2	Nascent and Future Applications	59
8	Summary and Future Outlook	59

1. Introduction

Positron Emission Particle Tracking (PEPT) is a nuclear imaging technique widely used in a range of scientific disciplines, including physics [1–4], chemical, mechanical, nuclear and civil engineering [5–10], and more recently the biomedical field [11–13]. PEPT allows the fully-three-dimensional tracking of one or more ‘*tracer particles*’ as they move through a given system. By using highly-penetrating 511 keV gamma rays, the interior dynamics of even large, dense, optically opaque systems may be viewed with high temporal and spatial resolution [14]. PEPT was originally developed at the University of Birmingham in the late 1980s [15, 16] and, despite extensive use in diverse applications, for the first three decades of its existence operated using largely the same fundamental algorithms originally implemented. The last decade, however, has seen an upsurge in the development of novel approaches to PEPT, including algorithms utilising cutting-edge machine-learning approaches [17, 18] and highly-efficient trajectory reconstruction algorithms with potential applications in the study of bloodflow, cancer metastasis and other biomedical applications [11]. Recent years have also seen the advent of the numerical modelling and simulation of PEPT systems [17, 19–21], a new capability with the potential to revolutionise both the development of PEPT systems and algorithms, as well as their application to scientific and industrial processes.

CONTENTS

4

While there exist multiple distinct PEPT methodologies, the underlying algorithms have been developed, tested, and demonstrated using vastly differing hardware. As such, it is near-impossible to draw direct comparison between these methods simply by analysing and comparing existing published work [17,22] – a problem we aim to address definitively in this paper.

In the following pages, we begin with an accessible introduction to and overview of PEPT, its fundamental methodology, and its grounding in nuclear physics (Section 2). We then, in Section 3, discuss both the underlying physics and practicalities of the detectors and radiotracers required to perform PEPT, and subsequently (Section 4) explain how this hardware, and indeed full PEPT experiments, may be numerically modelled [23,24]. In Section 5, we review all recently-developed and/or widely-used PEPT methodologies, providing links to interactive online resources through which the user may gain direct insight into the underlying algorithms thereof. In Section 6 we provide a first direct, quantitative comparison of the capabilities of different PEPT methodologies, utilising new results from a carefully-designed testing framework (Section 6.1). In Section 7, we provide a brief overview of the diverse applications (and potential future applications) of PEPT in science, industry and medicine, before summarising in Section 8.

2. The Fundamentals of PEPT

Before discussing PEPT and its underlying algorithms in detail, it is important to first ensure solid understanding of the fundamental principles upon which PEPT operates. In this section, we will present a basic explanation of the PEPT technique in an idealised system, before going on to introduce some of the major complicating factors which necessitate the use of advanced algorithms such as those explored in this paper.




Figure 1. Simple schematic illustrating (a) the emission of back-to-back γ photons from a radioactively-labelled tracer particle and (b) the detection of the emitted photons by a dual-headed gamma camera.

CONTENTS

5

2.1. PEPT in an Ideal World

Consider a single tracer particle, radioactively ‘labelled’ with a positron-emitting radioisotope – that is to say that a number of β^+ -emitting atoms have either been adsorbed onto, or created within, the tracer material[‡]. As positrons usually possess short mean free paths, λ (typically of the order of millimetres), especially in relatively dense media [25–27], one may reasonably expect that an emitted positron will annihilate with an electron either in or near the tracer itself. For the idealised case presented here, however, let us make the simplifying assumption that all positrons are emitted from the geometric centre of our particle, and annihilate immediately upon emission (i.e. $\lambda = 0$). The annihilation of a positron and electron will produce a pair of 511 keV γ rays with trajectories separated by $180 \pm 0.5^\circ$ [§] – i.e. they can be expected, to a reasonable approximation, to trace a straight line through the centre of the tracer volume (Fig. 1(a)).

Consider next that our tracer particle is placed within the field of view of a suitable detector system (a ‘positron camera’) capable of recording both the position, \vec{r}_i , and time, t_i , at which an incident photon i interacts with said detector system. A simple schematic of such a set-up is provided in Fig. 1(b), depicting a detector system known as a dual-headed planar camera, comprising a pair of flat, parallel detector heads. For such a system, if any two photons are registered by the opposing detector heads within a given resolving time, t_r – typically of the order of nanoseconds or hundreds of picoseconds [29] – it is assumed that both photons must have been emitted from the same positron-electron annihilation event. Assuming also, as above, that such photon pairs follow collinear trajectories and originate from the tracer particle, we can use the recorded values of \vec{r}_1 and \vec{r}_2 – the photons’ detected positions – to reconstruct the linear path taken, known as a ‘line of response’ (LoR). If several pairs of photons are emitted in quick succession, one may – for an ideal system – determine the position of the tracer by calculating the point of intersection of these lines, and thus the tracer’s position, allowing us to effectively track its motion, and thus extract also additional parameters, such as its pseudo-instantaneous velocity [22]. If a tracer’s trajectory is recorded for a suitably long period, one may also use PEPT data to extract a wide range of physical quantities and fields, ranging from self-diffusion coefficients [30–32] and velocity distributions [33] to three-dimensional density distributions [34]. While the present work does not focus on these parameters and their determination from PEPT data, the interested reader may find details thereof in references [22, 30, 35–37].

[‡] The processes through which this may be achieved are described in detail in Section 3.2.

[§] The possibility of a slight acollinearity arises due to the fact that the positron and electron are not *perfectly* at rest before annihilation. While the motion of the positron is normally largely arrested before annihilation, the momentum associated with the electron’s orbital state may nonetheless exert a small influence on the resultant photons’ trajectories. It should be noted that the value of $\pm 0.5^\circ$ is only indicative, but nonetheless provides a reasonable characterisation of the extent of acollinearity to be expected [28].

CONTENTS




Figure 2. Schematic diagrams illustrating (a) the emission and detection of multiple pairs of photons and (b) the reconstruction of their trajectories, from which a reconstructed particle centroid may in turn be obtained.

2.2. PEPT in the Real World

In reality, determining a tracer’s position from a given ‘sample’ of LoRs is somewhat more complicated. Firstly, one must consider ‘*positron flight*’ (the distance a positron will travel between β^+ emission and annihilation) meaning one cannot realistically assume that all photons will be emitted from a single point in space at which all LoRs neatly intersect. The extent of positron flight depends on several factors – the tracer size and material, the density of the surrounding medium/media, and the initial positron energy [25]. This ‘*spread*’ is further compounded by the detector’s finite spatial resolution [21], meaning that recorded values of \vec{r}_i all carry an inherent error. As such, detecting a particle is more complex than simply finding a point of intersection; there do, however, exist various ways in which particle position data may be extracted from a sample of LoRs, as we will see in section 5.

One must also consider that not all LoRs correspond to ‘real’ annihilation events, and even those which do may not correspond to the true trajectory followed by the resultant photons. Recall from above that a positron camera records the positions of, and thus produces an LoR from, *any two photons* detected within the system’s resolving time – i.e. it is possible that a pair of entirely unrelated photons may happen to be detected within t_r – a ‘*false coincidence*’. Such coincidences may be created, for example, due to secondary radiation from the tracer, i.e. if the labelling isotope’s decay results in direct γ -emission as well as β^+ -emission (‘*associated coincidences*’); if two annihilation events occur in rapid succession but only one photon from each pair is detected and/or due to the detection of background radiation (‘*random coincidences*’); or if one or both photons belonging to an annihilation pair become scattered, thus deviating from the collinear trajectory assumed by the detection algorithms (‘*scattered coincidences*’). These possibilities are illustrated Fig. 3.

Clearly, false coincidences will detract from the accuracy with which we may locate our tracer particle, and in severe cases even result in entirely false reconstructed locations. To maximise accuracy, we must find an effective means through which to separate ‘*true coincidences*’ from false. As will become evident in the coming pages, this is a surprisingly complex task, for which a wide range of approaches have been

CONTENTS




Figure 3. Schematic diagrams illustrating the various forms of false coincidence which may occur during PEPT imaging, as described in the main text.

– and continue to be – developed. In Section 5 we provide an overview of currently available approaches, and in Section 6 we put them to the test.

3. PEPT Hardware

3.1. PEPT Detectors

3.1.1. An Overview of Detector Physics– In this section we provide a brief overview of relevant detector physics concepts. For brevity, we provide only key terms and concepts necessary to the comprehension of the rest of this article; the interested reader, however, may find further detail and contextualisation in Appendix 1.

As is clear from the preceding section, to perform PEPT effectively, we need to be able to detect single 511 keV annihilation photons at high efficiency and with high energy and spatial resolution, and be able to precisely detect coincidences. The total number, n_s , of individual annihilation photons (‘singles’) detected, known as the true ‘singles event rate’ is determined by the activity, A , of the tracer, the fraction, γ , of total decays that yield positrons, the ‘efficiency’ (the fraction of photons from a source that a given detector is expected, on average, to register), ε , of the detector and the fractional solid angle subtended by the detector, Ω :

$$n_s = A\gamma\varepsilon\Omega, \quad (1)$$

CONTENTS

8

With high-activity tracers the singles rate measured is reduced by ‘*deadtime*’, the period of time when a detector has recently registered a photon interaction and, whilst processing the current interaction, cannot register any subsequent interactions. By considering the effect of deadtime through a ‘*deadtime factor*’, f_{DT} , we can use the above expression for the *total* singles rate to provide an expression for the *measured* singles rate:

$$m_s = n_s f_{DT}. \quad (2)$$

For very high activity sources, the deadtime can be so great that the camera becomes ‘*blinded*’ by photons; as such, perhaps somewhat counter-intuitively, more activity is not always better in PEPT imaging.

As is clear from the above, deadtime parameters can be a critical limiting factor affecting the temporal and spatial resolution of a PEPT measurement. For this and many other reasons, different positron cameras offer strongly differing performance, as we will discuss in the subsequent section.

3.1.2. Conventional Detector Systems (Ring and Planar Scanners)– The original PEPT detector system was a Multi-Wire Proportional Chamber (MWPC) positron camera developed in the 1980s at the University of Birmingham, in collaboration with the Rutherford Appleton laboratory, Oxford, and the Royal Marsden general hospital, Surrey [38,39]. As this article concerns *recent* developments, however, we will not discuss the MWPC system in detail – though the interested reader may find further information in Appendix 2. Nonetheless, as the first ever PEPT system, it is worth mentioning briefly as a benchmark against which subsequent hardware can be measured. Using the precursor to the Birmingham algorithm [40] (see section 5.1), the MWPC could locate a tracer moving at 1 m/s to within approximately 5 mm, approximately 50 times per second.




Figure 4. (a) The University of Birmingham’s ADAC Forté scanner. (b) The University of Cape Town’s ADAC Vetex (foreground) and HR++ (background) scanners. (c) The HR++ ring scanner at the University of Cape Town.

In 1999, the MWPC was replaced with an ADAC *Forté* camera, pictured in

CONTENTS

9

Fig. 4(a); the University of Cape Town acquired a functionally similar system, the ADAC Vertex, in 2011 (Fig. 4(b)). Both cameras represent exemplars of dual-headed ‘*planar scanners*’; the Birmingham Forté is in fact one of two positron cameras used in the algorithm comparison tests presented in section 6. Each of the ADAC systems’ dual detector heads comprises a single, large area ($590 \times 470\text{mm}^2$) sodium iodide (NaI) scintillation crystal, which converts incident 511 keV gamma photons into a cascade of lower-energy photons. These crystals are optically coupled to an array of 55 photomultiplier tubes (PMTs), which detect the photons from the aforementioned cascade. The relative number of photons detected by different neighbouring PMTs can be used to determine the ‘*light centroid*’, and thus the approximate position at which the original gamma photon interacted with the detector [41]. The spatial resolution of each of the camera’s two detector heads at 511 keV – that is, the precision with which it can ascertain the position of a single interacting photon – is approximately 6 mm. The maximum total *coincidence* rate of the camera – that is, the number of annihilation pairs of photons that can be detected per unit time – is approximately 100–150 kHz (for reference, the maximum *singles* rate of this camera is of the order of 4 MHz per detector head).

In addition to planar scanners, the other commonly-used positron camera geometry – indeed the most commonly used for modern medical positron emission tomography (PET) – is the ‘*ring*’ geometry, in which multiple, smaller detector crystals are arranged in a cylinder surrounding the patient/system of interest (see e.g. Fig. 4(c)). These scanners typically cover a smaller field of view than planar scanners like the ADAC, meaning only relatively small-scale systems, or specific sections of larger apparatus, can be studied. Ring scanners do, however, typically offer high sensitivity, and thus can be operated at higher event rates and with greater spatial resolution.

The Positron Imaging Centre (PIC) at Birmingham first investigated the potential for PEPT with fixed ring scanners in 2006-2007 [42], utilising several different systems [43–45]. The Cape Town facility hosts the unique ECAT HR++ (Fig. 4(c)), which comprises six rings of HR+ block detectors instead of the standard four, extending its axial field of view. Birmingham and Tennessee have investigated PEPT utilising small animal PET scanners with $\sim 1\text{mm}$ spatial resolution, which operate on the same principle used by the larger conventional ring scanners, described below. These scanners offer even better sensitivity than ‘normal’ (human) ring scanners, but at the cost of a still smaller field of view (typically $\sim 100\text{ mm}$ diameter), restricting yet further the size of the system that can be investigated. At the University of Tennessee, PEPT experiments have been conducted [19, 46, 47] using such scanners, including a Concorde Microsystems microPET P4 [48] and a Siemens Inveon [49]; a newer scanner, the Mediso MultiScan LFER [50], is now being explored for PEPT deployment at Virginia Commonwealth University. This scanner is designed for primate imaging, and as such offers a slightly larger field of view than other animal scanners.

Unlike the previously-described, single-crystal ADAC, ring scanners use ‘*block detectors*’ [51] (see Fig. 5); though not necessary to the comprehension of this article,

CONTENTS

10




Figure 5. Examples of individual block detectors. Large numbers of such blocks can be assembled into ring scanners such as those shown in Fig. 4(c) or the modular positron cameras of Birmingham and Cape Town shown in Fig. 6.

and thus omitted from the main text, the interested reader may find details pertaining to block detectors, their structure, and their operation in Appendix 3.

Individual detector blocks are grouped into units called ‘*buckets*’ which typically host 4 to 12 blocks. For Bismuth Germanate Oxide (BGO) crystals used in contemporary PEPT scanners, a single bucket may achieve an event rate of 1.5–2.0 MHz [52]. As the total singles rate for a camera system is simply the sum of the individual bucket singles rates, a ring scanner (which comprises a number of such buckets) can easily achieve a total singles rate of 30 MHz; the HR++ system at Cape Town, for example, can provide singles rates > 70 MHz for suitably active particles.

3.1.3. Modular Cameras– In PEPT, unlike PET, uniformity of the field of view – i.e. a symmetric geometry such as that achieved in ring and planar cameras – is not required, meaning customised, ‘modular’ geometries can be designed to accommodate large and/or unusually-shaped apparatus [53–55]. The block-bucket approach described above is, at its heart, inherently modular, with buckets able to operate independently of one another. As such, to within reason, it is possible to arrange buckets in any arbitrary geometry.

This concept forms the basis of the modular PEPT camera: blocks sourced from conventional ex-clinical ring scanners are coupled to a custom data acquisition system and arranged in a suitable geometry for a given system, with Monte Carlo simulations (discussed in section 4) used to optimise the detector layout. The detector modules are rugged and transportable, enabling PEPT measurements to be made on apparatus that cannot easily be moved, or for infrastructure requiring specific associated equipment not available in the laboratory, for example. Figure 6 shows a modular camera being used in tandem with the ADAC Forte at Birmingham to study a relatively large-scale, two-section fluidised bed. Here, the modular camera was used to observe the upper section of the bed and the ADAC to observe the lower.

CONTENTS

11




Figure 6. The modular (top) and ADAC Forte (bottom) cameras being used to study a fluidised bed system larger than can be accommodated in any conventional scanner alone.

3.1.4. Evolution of PEPT Camera Capabilities– As is clear from the preceding sections, one of the most direct manners in which to improve the capabilities of the PEPT technique is to simply employ a positron camera with a greater data acquisition rate. Fig. 7 provides an illustration of the progression of data acquisition rates over generations of positron cameras. The abscissa shows the *total* (i.e. true + false) coincidence data rate acquired, limited by system deadtime, for each positron camera analysed. The ordinate displays the number of these total events which correspond to true events. The points of inflexion illustrate the maximum count rate of each system, with deviations from the linear caused by deadtime processes, as discussed above. The historical trend has been in pushing the points of inflexion towards the upper right corner as data acquisition systems improve and deadtime factors are reduced. One of the major aims of the field is to continue pushing this boundary.

3.2. PEPT Tracers

The progress of PEPT and our ability to exploit an ever wider range of applications is inherently associated not only with the advancement of detector technology (as discussed

CONTENTS

12




Figure 7. The progression of data acquisition rates over generations of positron camera.

in section 3.1) and algorithms (as discussed in section 5) but also the development of new radiotracers. As the standard PEPT radiolabelling process has been described in detail in previous publications [56, 57], we limit our focus here to recent advancements and a brief overview of the underpinning principles, terminology and concepts.

3.2.1. Fundamental Concepts The radiotracers employed in PEPT differ significantly from those used in PET; in particular, a typical PEPT radiotracer comprises a solid, macroscopic particle, as opposed to taking the form of a fluid. Unlike in most conventional, (microscopic) molecule-based radiotracers, the introduction of a radionuclide to such a particle usually has a negligible effect on its original structure due to its large size [59], allowing PEPT to be used non-invasively.

If a truly non-invasive measurement is desired, it is necessary that the tracer particle used matches the physical properties (size, density, frictional coefficient etc.) of those constituting the system to be imaged. In the case granular systems, this can be achieved simply by using choosing a tracer comprised of the same material as the bulk – i.e. if we are imaging a fluidised bed of sand particles, simply irradiate a sand particle! Though

|| Note that the term ‘microscopic’ as used here does not refer to particles of the micron scale, which can be used in PEPT imaging. Rather, the terms ‘microscopic’ and ‘macroscopic’ are here used to distinguish between particles on the atomic/molecular scale and larger ‘granular’ particles comprising large numbers of individual atoms or molecules, respectively [58].

CONTENTS

13

not always possible due to the limitations of current radiolabelling techniques [59], in a majority of cases particles truly identical to those used in the bulk system under investigation can be radiolabelled in at least one of the three main ways described in this section.

When attempting to assess which – if any – labelling technique is appropriate for a given material, one will typically evaluate three parameters: the radiochemical yield (RCY), radiochemical purity (RCP) and radiochemical stability (RCS). The RCY simply corresponds to the amount of radioactivity successfully ‘attached’ to the material with respect to the starting amount of radioactivity; the RCP measures the presence of other radioactive elements within a sample; the RCS quantifies the strength of the material-radionuclide bond. For PEPT, we require high values of all three parameters. Clearly, a high RCY is desirable in order to get high activity and thus a high event rate. A low RCP – i.e. the presence of other, non- β^+ decaying radioactive species – will add unwanted noise (i.e. false LoRs) to measurements and lower the maximum activity that can be used without incurring problems due to deadtime. A low RCS means the radionuclide may easily become detached from the material during measurement. This leaching of radioactive material can prove problematic in several ways: not only is it reducing the activity of the tracer itself, but it is also introducing background noise and, if the leaching is significant, may even lead to false locations if the leached material accumulates in a given region of the system – i.e. the PEPT algorithms track the ‘free’ radionuclide as well as – or, worse, *instead of* – the radiolabelled particle itself.

There exist three main strategies for the radiolabelling of PEPT tracers – direct activation, ion exchange and surface labelling (see Fig. 8) – each of which we will introduce briefly below.




Figure 8. Simple schematic illustration depicting the three main ways in which atoms of a positron-emitting radioisotope (represented by yellow particles carrying a black trefoil) can be integrated into a tracer either by converting an existing atom through irradiation with a suitable ion beam (direct activation, top left), exchanging a radioactively inert atom with a positron emitter (ion exchange, top right) or simply adhering a suitable radioisotope to the existing atoms (surface labelling, lower centre).

CONTENTS

14

3.2.2. *Direct Activation* In many ways the most straightforward irradiation technique, in this method activity is induced *directly* in the desired tracer particle by bombarding it with a high-energy beam of accelerated ions (typically protons or Helium-3 nuclei) in order to generate a suitable positron-emitting radionuclide from a stable element within the tracer material. One of the most commonly-employed routes is to irradiate a material containing natural oxygen (^{16}O) with a 33 MeV ^3He beam to produce ^{18}F from ^{16}O atoms via the nuclear reactions:



and



The popularity of this particular reaction stems from two main factors – firstly, oxygen atoms can be found in many materials that are widely used in industry, such as glass, sand, alumina and zirconia; secondly, by directly irradiating water (H_2O) one can then go on to label a wide array of other materials using the *indirect* activation techniques described below. Other commonly used isotopes include Cobalt-55, produced via the bombardment of iron atoms with protons, Copper-61, produced from Nickel using a beam of deuterons, and Gallium-66, produced by bombarding Nickel with protons, with a variety of other processes also possible [60–62].

As well as producing ^{18}F , the process described in equation (4) also produces a number of short-lived ‘*side radionuclides*’ such as ^{10}C (half-life 19.3 s), ^{12}N (half-life 11 ms), ^{27}Si (half-life 1.16 s), ^{29}P (half-life 4.1 s), and ^{26}Al (half-life 6.4 s) [60]. To ensure a high RCP – i.e. the absence of large amounts of ‘*secondary radiation*’ which will add noise to the reading without increasing the true event rate – the tracer produced must therefore be allowed to ‘cool down’ for a short period of time (~ 20 minutes) after the reaction until these unwanted radionuclides have largely decayed away. Similar procedures are also typically required for other radionuclides, with the duration of the cool down period varying according to the half lives of both the desired and unwanted radioisotopes.

One of the great advantages of direct activation is that, as the radionuclide is created from a stable element which is a part of the particle’s structure, detachment of the radionuclide is extremely unlikely – i.e. particles produced via direct activation typically have an excellent RCS.

Due to the significant heating effects inherently experienced when a particle is bombarded with high-energy ion beams, the direct activation process is typically limited to relatively large particles (> 1 mm) and materials with relatively high melting points. For particles which do not fulfil these criteria, alternative approaches – i.e. indirect activation – must be explored.

¶ Helium-4 nuclei or deuterons may also be used, but this is less common as fewer widely-used reactions require these ions.

CONTENTS

15

3.2.3. *Indirect Activation* There are two main methodologies by which indirect activation may be achieved: ion exchange and surface labelling. In the former, instead of a stable atom within the material of interest being *converted into* a positron emitter, an atom is instead removed from the material and *replaced with* a suitable positron emitter. This is typically performed by placing a suitable material such as an ion exchange resin [63] in a liquid containing a high concentration of free ^{18}F ions. Under suitable conditions [56], the free fluorine ions will preferentially attach to the resin, thereby completing the labelling process. The chemistry of this process is briefly detailed in Appendix 4. The main advantage of the ion exchange method is the relative simplicity of the radiolabelling process. A major disadvantage, however, is the weakness of the bond between the particle and the radionuclide, giving a low RCS, and thus a significant possibility of activity leaching into the surrounding system – especially in the case of measurements involving fluid flow.

The surface labelling strategy is based on the affinity of certain materials toward specific radionuclides. By exploiting different chemistries, through the careful choice of material and the radionuclide, it is possible to attach positron-emitting nuclei to the surface of a particle by strong interactions. In the field of nanomaterials, this method has been applied for the radiolabelling of different inorganic particles such as those made of hydroxyapatite, silica, or iron oxide with radionuclides such as ^{18}F , ^{68}Ga , ^{64}Cu or ^{89}Zr , among others [64].

The surface labelling strategy is currently being used by researchers at King's College London to facilitate the activation of silica particles as small as $1\ \mu\text{m}$ with ^{68}Ga for use in PEPT applications. The radionuclide, here ^{68}Ga , is able to react with the silanol groups of silica, forming strong bonds in a short period of time. A similar approach may be used to increase the affinity of ^{18}F toward various different materials, for example by activating the surface of the material of interest using Fe^{3+} salts previous to the incubation with ^{18}F . This strategy allows the successful radiolabelling of materials such as microcrystalline cellulose (MCC), quartz, calcite, apatite or coal particles.

Aside from the specific examples described above, there are many possibilities for the radiolabelling of different nanomaterials, and hence possible applications for PEPT. There are, however, several challenges associated with the labelling of nanoscale materials. These include a decrease in the potential for high levels of radioactivity per particle and in the isolation of a single particle. Future work to address these challenges is ongoing at King's College, with the aim of expanding the application of PEPT into new and exciting areas, perhaps most notably in the biomedical field, where the use of small, biocompatible tracers may prove revolutionary.

CONTENTS

16

4. Simulating PEPT

4.1. The Value of Simulation

In recent years researchers have begun using numerical modelling to simulate PEPT hardware and experiments, creating artificial PEPT data from simulated fluid flows [20] and even producing highly faithful digital twins of real PEPT systems [21]. For a technique such as PEPT, the availability of such models can be highly valuable for a variety of reasons. As evident from the previous section, to perform ‘real’, physical PEPT experiments, one requires a significant amount of expensive hardware and infrastructure. Moreover, even with the relevant facilities in place, the inherent material, staff and energy costs associated with PEPT imaging make it a cost-intensive process. In other words, when using PEPT, one wants to make as efficient use of one’s time as possible, and avoid unnecessary tests and/or failed experiments. By using simulations, it is possible to test the viability of an experiment, or the suitability for PEPT imaging of a given experimental system, using only a laptop, thus minimising the chance of wasted experimental time and maximising the benefits gained from PEPT experimentation. Simulations can also be used to maximise these benefits further by predicting in advance the optimal parameters (e.g. tracer activity, tracer number, head separation, system positioning, acquisition time...) with which to perform any given set of experiments and, as discussed in the preceding section, to help optimise the geometries of custom PEPT detector systems.

An additional major benefit is the provision of a valuable ‘sandbox’ in which to explore the fundamental physics of the technique, to develop new theoretical and empirical models, and – as is the case in the present work – develop and test algorithms.

Though a fairly recent development within the field, it is expected that in the coming years the use of simulation will become increasingly commonplace for PEPT research, as it already has with other nuclear physics techniques [65–68].

4.2. Modelling PEPT Hardware

As we have seen in the preceding sections, both radiation detectors (see section 3.1) and radiotracers (section 3.2) are central to PEPT experiments [19] [69]. As such, if we wish to create an accurate numerical simulation of a PEPT experiment, it is vital that we are able to faithfully model positron sources, radiation transport through experimental geometries, and synthetic pulse-processing chains, amongst other crucial aspects of the emission and detection processes. The Geant4 Application for Tomographic Emission (GATE) software [70, 71] provides an invaluable tool for such modelling. In GATE, geometries can be created manually from design drawings or imported as detailed mesh files created using suitable CAD software. Particle interactions with the imported or created geometries are passed through a series of conditions emulating how the electronic pulse-processing systems of ‘real’ positron cameras would respond [72]. This step, called the digitizer, is unique to each detector system. The digitizer is also partially responsible




Figure 9. *Left:* A simulated fluidised bed can be studied in a GATE simulation by using the particle volume fraction at each timestep to create materials and prescribing the extracted positions of a single particle as a positron emitting source. *Right:* A CT scan of a chest phantom voxelised inside a GATE simulation.

for recreating a crucial aspect of each detector system – the coincidence count-rate, and its variation with source activity – and can also include the effects of deadtime [73], as touched upon in section 3.1.

As with any numerical simulation, it is vital that GATE models are validated against experimental data [74]. To validate the GATE model of a detector system, a common set of experiments using the real and GATE modelled detectors should be conducted, such that the values attained numerically and experimentally can be compared, and thus provide a quantitative assessment of the model’s accuracy. The National Electronics Manufacturers Association (NEMA) provides a standardised protocol, NEMA NU-2, for testing the performance of PET scanners; this framework has been used in validating the two GATE models used in section 6 of this work: the Siemens ECAT EXACT HR+ ring scanner and the Phillips ADAC Forte dual-headed positron camera [75] [21] [76]. If values obtained from the virtual and experimental NEMA tests agree to within an acceptable tolerance, then one may assume that said model provides a reasonable simulation of the real detector system. Examples of such validation processes can be seen in [21].

The NEMA-validated GATE models of the ADAC and ECAT systems used in the present work can be freely downloaded from open-source repositories linked, respectively, [here](#) and [here](#).

4.3. Simulating a PEPT Experiment

As mentioned above, the use of a validated GATE model to study PEPT experiments requires only a CAD model of the experimental geometry itself and relevant details (e.g.

CONTENTS

18

size, material and motion) of the relevant tracer(s) to be tracked. Tracer movement can be prescribed analytically by assigning a source velocity vector; alternatively, discrete movements at specific timesteps can be read in from another file, allowing trajectories produced via, for example, the Discrete Element Method [77] or other numerical simulations to be directly converted into ‘pseudo-PEPT’ data. When a GATE simulation is conducted using realistic source movement, the detected LoRs can be used to reconstruct the tracer trajectory as in a ‘real’ PEPT experiment. This pseudo-PEPT trajectory can then be compared to the known ‘ground truth’ data – i.e. the originally prescribed positions – giving the user an estimate of the accuracy of a given PEPT experiment prior to experimentation. Further, the effect of the detector settings, source activity, and experimental conditions can be altered in the GATE simulation to determine optimal conditions for the ‘real’ experiment. Finally, as we will see in section 6, such a comparison can also allow us to test and compare the capabilities of different PEPT algorithms.

As well as source movement, the movement of the system being imaged can be included in GATE simulations. In multiphase systems such as fluidised beds, large transient voids created by bubbles of gas moving up through the dense medium can significantly affect the scattering of gamma rays by reducing the amount of material which must be traversed by the rays before they are detected. If a DEM simulation of the bed is created, the particle volume fraction in discrete cells can be prescribed into the GATE simulation and updated during each timestep. This process, known as ‘voxelisation’ can also be used to create detailed representations of the geometry of complex systems where the internal structure is not known by using a Computed Tomography (CT) scan [78]. CT images can be placed inside a GATE simulation and prescribed material properties such that the scattering of gamma rays inside a complex system are approximated [79].

5. PEPT Methodologies – an Introduction

In the preceding sections, we have provided an introduction to the fundamental principle of PEPT, and the tracers and detector systems through which it operates. At the end of Section 2, we discussed the complexity of extracting accurate particle positions from the mix of true and false lines of response provided by raw PEPT data, and accounting for the inherently limited resolution of PEPT detectors and spread caused by positron flight. In section 3, meanwhile, we discussed the push toward faster and more accurate detection through the development of improved hardware, and the development of tracers capable of carrying higher levels of β^+ activity. In this section, we provide a detailed overview of the various methodologies developed to address the challenges highlighted in Section 2, and push the boundaries of PEPT’s spatial and temporal resolution further than hardware and tracer development alone. Considering the aims of this article, we place a particular focus on the most recent and most promising algorithms, and those which are most widely used. In each case, we will provide an overview of the methodology

CONTENTS

19

and underlying algorithms, and briefly summarise the major strengths and weaknesses thereof. A more detailed, quantitative analysis of these will then be given in section 6.

5.1. The Birmingham Method

The Birmingham Method, developed at the University of Birmingham by the original developers of the PEPT technique, David Parker and Michael Hawkesworth, is the oldest yet still, to date, the most widely-used PEPT algorithm. The fundamental principle of the method is, for a given set of LoRs, to discard false events iteratively until only true LoRs remain, finding an accurate particle centroid from the remaining data [80, 81].




Figure 10. (a) Simple visualisation of the perpendicular distance, δ , between an LoR and an arbitrary point (x, y, z) in space, as discussed in the main text. (b) For a pair of perfectly-intersecting photons emitted from a tracer, the minimum distance point, \vec{r}_{MDP} , will correspond exactly to the point of intersection. (c) The presence of false coincidence events may pull the MDP away from the true particle centre.

The algorithm considers a sample containing a user-defined number, N_s , of time-consecutive LoRs from a full PEPT data set. Each LoR is defined by its start- and end-points (i.e. the positions at which the recorded photons interact with the detector) which can be given in Cartesian coordinates as (x_1, y_1, z_1) and (x_2, y_2, z_2) (see Fig. 10). The smallest distance, $\delta(x, y, z)$, between the LoR defined by these points and any given point, (x, y, z) , in three-dimensional space, can be determined as:

$$\delta^2 r^2 = [(x - x_2)z_{12} - (z - z_2)x_{12}]^2 + [(y - y_2)x_{12} - (x - x_2)y_{12}]^2 + [(z - z_2)y_{12} - (y - y_2)z_{12}]^2 \quad (5)$$

where $\chi_{12} = \chi_1 - \chi_2$ ($\chi = x, y, z$), and $r = \sqrt{x_{12}^2 + y_{12}^2 + z_{12}^2}$ is thus the scalar length of the current LoR of interest. For a given sample of N_s LoRs, one can determine the root mean squared distance, $D(x, y, z)$, of point (x, y, z) from *all* LoRs in our sample as:

$$D = \sqrt{\frac{1}{N_s} \sum_i^{N_s} \delta_i^2} \quad (6)$$

CONTENTS

20

By finding the minimum of the above function, which may be achieved by simultaneously solving:

$$\frac{\partial D}{\partial x} = 0, \quad \frac{\partial D}{\partial y} = 0, \quad \frac{\partial D}{\partial z} = 0, \quad (7)$$

we can determine the point in space which minimises the total distance to all LoRs, known as the ‘*minimum distance point*’ (MDP). As the MDP will be influenced by both true *and false* LoRs, if we wish to determine the position of our tracer accurately, we must first eliminate erroneous data. This is achieved using a simple iterative process: for an initial sample of N_s LoRs, an initial MDP is determined as outlined above, and all LoRs whose perpendicular distance, δ_i , from this point are more than 20% greater than the RMS distance D (i.e. $\frac{\delta_i}{D} > 1.2$) are discarded as ‘*corrupt*’. A new MDP is then determined from the remaining LoRs, and the process repeated until only a user-defined fraction f of the initial sample remain. A simplified example of this process is shown in Fig. 11; in reality, a significantly larger number of LoRs (typically of the order of 100) is normally used.




Figure 11. Simplified schematic illustrating the operation of the Birmingham method for a sample containing an initial $N_s = 15$ LoRs, and processed using an f value of $\frac{1}{3}$. In each panel, dashed blue lines represent recorded LoRs, while black crosses represent the minimum distance point (defined in the main text) of this sample of LoRs. Panels (a), (b) and (c) correspond, respectively, to the initial sample of LoRs, an intermediate step in the algorithm, and the final remaining LoRs used to calculate the particle’s true centroid (see main text).

The Birmingham method may be relatively simply extended to the case of multiple tracers [82,83]. For a given sample of LoRs containing N_t tracers, the algorithm begins as before, iteratively determining the centroid position of a single particle. In the multiple particle case, it is typically beneficial if the iterative discarding of LoRs is a little more gradual; as such, at each step of the iteration only LoRs whose distance from the current MDP are 50% greater than the mean are discarded. Once this first tracer has been located, the process is repeated using the same sample, but omitting all LoRs identified as belonging to the original tracer. This process is repeated until all tracers have been located. Clearly, for each additional particle located, the initial pool of LoRs used to determine its location decreases in size, leading to inherently reduced statistics – that is to say, the more tracers introduced into the system, the lower the precision with which they may be detected. This approach has been shown to track up to four particles

CONTENTS

21

successfully, with a minimum separation of approximately twice the detector's inherent spatial resolution.

An interactive Python notebook illustrating the key steps in the Birmingham algorithm may be found [here](#). A complete implementation of the Birmingham method may also be found on Positron Imaging Centre github repository, <https://github.com/uob-positron-imaging-centre/>.

Strengths	Weaknesses
Arbitrarily high accuracy for stationary tracer	Can only track limited number ($\lesssim 4$) of individual tracers
Relatively computationally efficient	Requires <i>a priori</i> knowledge of N_t
Mesh-free method	Accuracy decreases sharply with N_t

Table 1. Major strengths and weaknesses of the Birmingham method.

5.2. The Line-Density Method

The line-density method [84] was developed in 2012 at the University of Cape Town. Though no longer used on its own, it forms the foundation of several subsequently-developed algorithms discussed and tested later in this paper, and as such will be addressed in some detail.




Figure 12. Simplified 2D image representing the voxelisation of the computational volume and the calculation of line density. In this image, cells intersected by lower (higher) numbers of LoRs are indicated by darker (brighter) shading, allowing us to easily discern the voxel containing the peak line density and thus the approximate tracer position.

CONTENTS

22

The method begins by discretising the computational volume (i.e. the camera’s field of view) into a series of three-dimensional voxels, as illustrated in Fig. 12. A numerical value is then assigned to each of these voxels corresponding to the total number of LoRs passing through it. Clearly, the closer to the source (i.e. tracer) a voxel lies, the higher the number of lines likely to pass through it. An estimation of the tracer location may therefore be determined by finding the voxel containing the highest line density. This estimate may be improved and a sub-voxel resolution achieved by considering a one-dimensional ‘row’ of voxels passing through this maximum along each Cartesian axis (x, y, z). For each of these 1D rows, one may plot the distribution of counts along the relevant axis.

If the tracer’s dimension is smaller than the pixel size, the distributions obtained may then be fitted with Gaussian functions, whose means and widths can be taken, respectively, as the location and error therein along each axis.

To track multiple particles, the approach outlined above is again applied, but this time looking for *local* maxima in the line-density matrix. To do so, both the number of tracers to be tracked and their approximate initial positions within the system must be specified when initiating the tracking process. Each particle is assigned a specific ‘region’ based on a visual inspection of the first density matrix. For each subsequent timestep, the regions in which to search for each given particle are then calculated based on the respective tracer’s previous position and velocity, allowing a reasonable estimation of its expected position for the following timestep. However, if the regions assigned to two or more distinct tracers overlap, the algorithm will record the position of only one, typically the more active. One must be careful, therefore, when deciding on the size of the search – too big, and tracers’ regions will often overlap; too small and a quickly and/or erratically-moving tracer may fall outside its predicted region. To address this latter possibility, the search region’s dimension, d , is typically constrained to satisfy the inequality $d > |a|(\delta t)^2$, where a is the particle’s maximum acceleration.

Under suitable conditions, the line-density method is capable of achieving a similar level of accuracy to the Birmingham algorithm. However, the use of a discretised mesh introduces certain problems not encountered in the original method, notably the grid-size dependence of results (thus requiring additional optimisation and sensitivity analysis), and potential for pixelisation effects for slower-moving tracers.

An interactive Python notebook illustrating the key steps in the Line Density algorithm may be found [here](#).

5.3. G-Means Clustering

The G-means method of Wiggins et al. [85] was developed at the University of Tennessee, Knoxville to enable multiple particle tracking using an arbitrary number of tracers. This method begins by creating a line density matrix as described in section 5.2. First, a check is made to ensure that the maximum line-crossing value is above a preset threshold value. If so, the grid of voxels is first high-pass filtered based on a fraction (usually 0.25-

Strengths	Weaknesses
Relatively computationally efficient Relatively simple to implement	Requires <i>a priori</i> knowledge of N_t Requires <i>a priori</i> knowledge of initial tracer positions Grid-size dependent results Subject to pixelation for slower-moving tracers

Table 2. Major strengths and weaknesses of the line density method.

0.5) of the peak line-crossing value. The remaining voxels are then clustered via G-means clustering [86–88] – modified k -means [89] algorithm – to determine their positions.

It is worth noting that the G-means algorithm was in fact not the first PEPT algorithm to make use of clustering – this was the K-Medoids method, developed by Gundogdu [90–92]; this method, however, was never widely adopted, and hence is not discussed in the main text. The interested reader, however, may find full details in Appendix 5.

Particles are identified as follows: Firstly, the LoR-crossing data are treated as points in \mathbb{R}^3 positioned at the centre of each voxel with a multiplicity corresponding to the number of line crossings at each point. Clustering begins with a $k=1$ clustering (i.e. a universal centroid calculation) and subsequently splits or accepts clusters based on a statistical test. Thus it grows the number of clusters until the natural number of clusters in the dataset is reached, avoiding the potential false convergence errors of k -means that can be caused by a poor initialization. The decision to split or accept a cluster is based on its adherence to a Gaussian fit.

For splitting, we consider a set \mathcal{S} of points with data already sorted into k clusters, $X_j \subset \mathcal{S}$, with $j \in \mathbf{J} = \{1, 2, \dots, k\}$. For a specific cluster, $X_m, m \in \mathbf{J}$, containing n points, the covariance matrix is calculated, and power iteration used to identify its main principal component and corresponding eigenvalue λ . Two daughter centroids, \mathbf{c}_1 and \mathbf{c}_2 , are initialized along the main principal component, a distance $\pm\sqrt{2\lambda/\pi}$ from the centroid of X_m , and k -means clustering with $k = 2$ used to re-cluster the data with \mathbf{c}_1 and \mathbf{c}_2 used as initial cluster centres. After converging on two new centroids, \mathbf{c}'_1 and \mathbf{c}'_2 , the data in X_m are projected onto the line $\mathbf{v} = \mathbf{c}'_1 - \mathbf{c}'_2$, such that $x'_i = (\mathbf{x}_i \cdot \mathbf{v})/|\mathbf{v}|$. This new dataset, X'_m , is a 1-D representation of X_m , normalized to have mean 0 and variance 1.

A 1-D Anderson-Darling (A-D) test [93] is used to determine if X'_m is normally distributed. For each of the n values $x'_i \in X'_m$, let $z_i = F(x'_i)$, where F is the $N(0, 1)$ cumulative distribution function, the A-D statistic, A^2 , is first calculated as

$$A^2(Z) = \frac{1}{n} \sum_{i=1}^n i = 1n(2i - 1)[\ln(z_i) + \ln(1 - z_{n+1-i})] - n. \quad (8)$$

Then, as the mean and variance must be estimated from the measured data, a correction

CONTENTS

24

is applied to A^2 as

$$A_*^2(Z) = A^2(Z) \left(1 + \frac{4}{n} - \frac{25}{n^2} \right). \quad (9)$$

If the value of A_*^2 is below a user-defined critical value, the cluster X_m is considered sufficiently Gaussian and accepted. If it is not, it is split, and another clustering performed with $k' = k + 1$. This process is continued until all clusters pass the A-D test. When all clusters pass the A-D test, the calculated centroids are taken as tracer positions. Fig. 14 shows a visual representation for the case of three activated particles [85].




Figure 13. Visual depiction of clustering via G-means. Clustering begins by treating the data as one cluster (a). This is then split into two clusters (b) and three (c) along the principle axes of each cluster. Further splitting is rejected at this point, it being assumed that the algorithm has found the total number of ‘natural’ clusters (i.e. the locations of all tracers within the field of view).

After particles are identified for all timesteps, they are linked into continuous trajectories, typically using the modified nearest-neighbours algorithm described by Wiggins et al. [85]. If between-frame displacements are particularly large (i.e. when imaging rapid flows and/or with low data acquisition rates), a more robust linking algorithm, described in [47, 94], may be used. After this linking process, a second ‘re-linking’ step is performed in which the start and end points of all trajectories are considered. If the beginning of one trajectory and the end of another are found to be very close in time and space, the apparent gap is considered to be an instance of occlusion, and the two trajectories are merged [94]. Statistical noise in the reconstructed trajectories may be subsequently removed by convolution with Gaussian kernel, as described further in section 5.5.

The G-means reconstruction algorithm has been shown to yield comparable accuracy to the Birmingham algorithm for single-particle location [85], and can track up to 17 particles simultaneously. A major advantage of this method is that it facilitates

CONTENTS

25

the tracking of multiple tracers without *a priori* information on the number or initial positions thereof. It also allows particles to enter and leave the field of view, enabling experiments in test sections with recirculating flow loops. However, for large N_t the clustering routine employed in this method is computationally expensive and shows poor convergence. Specifically, the selection of the A-D critical value (used to accept/reject the splitting of clusters) therein is based on user experience and is not well-founded in statistics due to the discrete nature of the data used. The number of clusters found in each timestep is very sensitive to the selection of this value. As such, the algorithm will often converge on an incorrect number of particles when many are present if the A-D critical value is not chosen properly. This can result in both false-positive and false-negative detections, which must be accounted for during trajectory linking. Furthermore, as this method uses discrete grid elements for line density tallying, there is some sensitivity to the selection of the grid size. In spite of these challenges, the utility of this method for flow measurements with multiple tracers has been demonstrated by Langford et al. [19].

An interactive Python notebook illustrating the key steps in the G-means algorithm may be found [here](#).

Strengths	Weaknesses
Capable of tracking up to 17 tracers Does not require <i>a priori</i> knowledge of N_t	Computationally expensive for large N_t Sensitivity to A-D critical value Grid-size dependent results

Table 3. Major strengths and weaknesses of the G-means method.

5.4. Spatiotemporal B-spline reconstruction (SBSR)

Spatiotemporal B-spline reconstruction [11], a technique developed in 2014 at Stanford University, aims, in brief, to find the spatiotemporal trajectory that is most consistent with the recorded LoR data by minimizing an objective function. Unlike prior PEPT the SBSR method's intended purpose is focused predominantly on the biomedical sector, having recently been demonstrated tracking the migration of single cells in live subjects [95]. To this end, the algorithm is designed so as to allow objects to be tracked using a minimal number of LoRs.

The trajectory of a tracer is modelled as a three-dimensional B-Spline [96] of the continuous temporal variable t . Each of the three spatial coordinates is modelled as a weighted sum of spline basis functions, $B_i(t)$; therefore, each trajectory, $\vec{r}(t)$, is represented by three sets of basis coefficients, $\vec{a}_i = (a_i^x, a_i^y, a_i^z)$. That is to say

$$\vec{r}(t) = \sum_{i=1}^N \vec{a}_i B_i(t). \quad (10)$$

CONTENTS

26

The number, N , of basis functions can be adjusted by the user for each dataset, considering the number of counts available and the presumed tracer velocity.




Figure 14. Simple schematic illustrating the SBSR method. The algorithm uses an least-squares approach to determine the best-fitting three-dimensional spline for a given set of points – i.e. the single trajectory which passes, on average, closest to all LoRs within a given sample.

With the basis functions set-up, a convex optimisation algorithm is used to find the spline weights that best match the data. The objective function used is the sum of the squared distances between each recorded coincidence event and the time-matched location along the estimated spline trajectory. A decrease in the objective function means that the reconstructed spline trajectory lies closer to the recorded events. As the objective function is convex, the optimisation is guaranteed to converge to a unique trajectory.

As the objective function's distance penalty is squared, false coincidences can have an outsized impact on the estimated tracer position. Therefore, the distance penalty is capped beyond a maximum threshold, which is set according to the spatial resolution of the PET scanner⁺.

5.5. The Feature Point Identification Method (FPI)

The feature point identification (FPI) method [97] was developed at the University of Tennessee in 2011. The method begins with the creation of a line-density grid, as described in section 5.2. As can be seen in Fig. 12, the grids created can in many ways

⁺ As the SBSR method is under patent, we do not provide here an interactive example as for other techniques.

Strengths	Weaknesses
Can track particles with very low activity Mesh-free method	Currently only tracks single particles

Table 4. Major strengths and weaknesses of the SBSR method.

be considered analogous to a digital image, with each cell corresponding to a pixel with an intensity representing the line density of (i.e. total number of pixels crossing) said pixel. The FPI method exploits this analogy by using techniques are borrowed from image processing and optical particle tracking [98,99] to identify, locate and track tracer particles.

The ‘image’ corresponding to the line-density grid created by a given sample of LoRs is first smoothed via a convolution-based filter. This is typically performed with a box-car kernel over a cubic region of width $2f+1$, where f is the smoothing size (usually 1 voxel). If $N^t(x, y, z)$ represents the number of LoR-crossings at voxel position (x, y, z) during timestep t , then the smoothed image can be computed as:

$$N^t(x, y, z) = \frac{1}{(2f+1)^3} \sum_{i=x-f}^{x+f} \sum_{j=y-f}^{y+f} \sum_{k=z-f}^{z+f} N^t(i, j, k). \quad (11)$$

For optical particle tracking, Sbalzarini and Koumoutsakos [98] used convolution with a Gaussian kernel for image smoothing. This can also be employed for PEPT in cases in which significant background noise is expected. Examples of the raw line-crossing grid and the smoothed image can be seen in Fig. 15.




Figure 15. Visual depiction of FPI processing, showing unfiltered line-crossing grid (left) and filtered image (right). Final tracer positions are indicated with \times 's. Experimental data shown are the same as those seen in Fig. 14.

Having smoothed the image, approximate particle positions can then be estimated

CONTENTS

28

as local intensity maxima in N^t . These local maxima are taken simply as the voxels possessing intensities in the upper r^{th} percentile of a given time frame with a line-crossing value greater than any other voxel within a cube of side-width $2w + 1$. The value of r can be adjusted based on the amount of noise (i.e. the preponderance of false LoRs) in a given image and the relative ‘brightness’ (i.e. detected activity) of each particle.

The parameter w serves as an apparent particle image size, where the image of each particle is a cube of high LoR density centred at the maxima. However, w also inherently places restrictions on the between-particle separation distances for which individual particles can be successfully resolved in a given timestep. If two or more particles are separated by less than w voxels, the particle that appears brightest (i.e. whose peak-intensity pixel has the highest line-density) will be detected, but the other(s) will be ignored; worse, the presence of the other (non-detected) tracer(s) will bias the recorded position of the detected particle toward the centroid of all tracers involved. Thus, if two or more particles approach closely, only that having the highest activity is likely to be followed. The position history of the other particle will therefore be split into two trajectories – another factor which may cause issues in the post-processing of the data collected, especially for cases in which it is important to assign specific trajectories to specific particles. Based on these significant potential complications, it is clear that – as far as possible – the value of w should be chosen such that it is smaller than the expected inter-particle separations for a given experiment.

Following the identification of local maxima as described above, final tracer positions are calculated within each image. Current implementations of the FPI method [46] calculate tracer locations using 1-D Gaussian fits in each spatial direction, centred at each local maximum, imitating optical peak location routines [100]. The position of each particle is taken to be the mean of the fitted Gaussian, and the uncertainty is calculated based on the standard deviation σ as $\Delta x = \sigma / \sqrt{N_{LoR}}$. N_{LoR} can be estimated for a given particle as the number of LoR passing within $(w+0.5)$ voxels of its calculated position. Tracer locations from individual timesteps can then be linked into trajectories by the nearest-neighbours method described in [85] and section 5.3, or the projection method described in [47, 94].

After trajectory reconstruction, measurement noise is reduced through the convolution (in time) of each trajectory with a Gaussian kernel of the form

$$k(\tau) = A \exp\left(-\frac{\tau^2}{\sigma^2}\right), \quad (12)$$

as used in optical particle tracking studies [101]. Here, σ is the effective filter size (in timesteps) and should be selected to be smaller than the smallest time scale of interest in the system interrogated. By convolving each trajectory with differentiated kernels, instantaneous velocity and acceleration can also be calculated for each tracer [97].

The multiple particle capability of FPI has been demonstrated for reconstructions of up to 100 particles via simulation [97] and 80 particles in an experiment [94]. Its primary drawback is the aforementioned inter-particle resolution limit set by the parameter w . In

CONTENTS

29

most measurements, this manifests as a temporary loss of information for some tracers but, in general, has little effect on overall measurement outcomes. The FPI method, like the other grid-based methods previously described, also displays a degree of grid-size dependence; however, the aforementioned use of Gaussian fitting during particle localization has been seen to remove any significant biasing due to pixelation.

An interactive Python notebook illustrating the key steps in the FPI algorithm may be found [here](#). A complete implementation of the FPI method may also be found on Positron Imaging Centre github repository, <https://github.com/uob-positron-imaging-centre/>.

Strengths	Weaknesses
Can detect at least 100 tracers Does not require <i>a priori</i> knowledge of number of tracers	Particle spacing limited by parameter w Grid-size dependent results

Table 5. Major strengths and weaknesses of the FPI method.

5.6. Voronoi-Based Multiple Particle Tracking (VMPT)

The VMPT method [102], developed by Blakemore, Govender *et al.* at the University of KwaZulu-Natal, unlike preceding techniques, uses three-dimensional Voronoi tessellation [103–105] to create a detailed, spatically-non-uniform line-density map, with statistical filtering [106] and DBSCAN clustering [107, 108] techniques used to determine particle positions therefrom. The trajectories are then extracted using multiple target tracking (MTT) [109–112], specifically the multiple hypothesis method [113].

Like the K-Medoids method briefly mentioned above, the VMPT technique has not, to date, been widely adopted, and has since been superseded by the Spherical Density Method, described below. As such, for brevity, its details are not discussed in the main text. Nonetheless, the interested reader may find full details in Appendix 6.

5.7. The Spherical Density Method

The spherical density method [114] of Odo, Govender *et al.* represents a marked departure from prior methodologies in that it requires a significant amount of *a priori* information regarding the tracer particles used and the distribution of beta-emitting material thereupon.

The method's operation is predicated on two main heuristic assumptions:

- I. A positron's mean free path, λ , between emission and annihilation is smaller than the tracer particle radius, R , meaning that one may reasonably expect a higher concentration of annihilation events inside the tracer medium than outside*

* For the case of a tracer in which the radionuclide is applied as a surface coating (Fig. 16), the

CONTENTS

30




Figure 16. Simple schematic illustrating the two types of radiolabelled spherical particle considered in the spherical density method. Panel (a) represents the case of a positron source placed at the centre of a dense spherical tracer, and panel (b) a similar tracer instead coated with a positron emitter. In both cases, red crosses represent positron-electron annihilation centres. The distances R and r^* represent, respectively, the radius of the tracer and the radial distance across which annihilation centres are expected to be most densely and evenly distributed [114].

- II. The highest and most uniform concentration of annihilation centres lies within a radial distance r^* from the radial centre of the positron-emitter with which the tracer is labelled, with $r^* < R$. The two main cases considered in the original work of Odo *et al.* [114] – a particle with a radiolabelled central core, and one with a surface coating of positron-emitting material – are depicted schematically in Fig. 17.

These simple assumptions can be used as a robust means of identifying and thus removing corrupt LoRs from a data set. Let us consider a simplified, two dimensional data set, pictured in Fig. 17. From visual observation, one can intuitively determine that lines 1, 4 and 9 are likely to be corrupt. Beyond this, however, it is harder to distinguish corrupt from true LoRs.

To sort true from false LoRs, the SDM considers the shortest distance, d_{i-j} between each given pair of LoRs,

$$d_{i-j} = \frac{|\vec{P}_i \vec{P}_j \cdot \vec{n}|}{\|\vec{n}\|}, \quad (13)$$

with \vec{P}_i and \vec{P}_j representing the closest points on a given pair of LoRs l_i and l_j and \vec{n} the unit vector connecting them. According to assumption II. above, for any uncorrupted LoR l_m passing through our tracer, any other uncorrupted LoRs corresponding to the same tracer must satisfy the constraint $d_{m-n} \leq R$ ($m \neq n, n \in [1, N]$). As such, the tracer position can be expected to be represented by a ‘set’, s_i , of multiple LoRs all satisfying this condition.

additional assumption that the tracer medium is denser than the surrounding medium must also be considered.

CONTENTS




Figure 17. *Left:* Hypothetical example of a small sample of LoRs in a two-dimensional plane. *Right:* Table detailing the total number of interacting LoRs (i.e. cases in which an LoR passes within a distance $d_{i-j} \leq R$ of the current LoR of interest) belonging to the sample, s_i , pertaining to each given LoR, l_i .

Let us consider, as an example, the set, s_5 belonging to LoR l_5 from Fig. 17. We can see that l_5 possesses a total of 6 points of nearest approach falling within the constraint $d_{5-n} \leq R$. To determine which of this ‘set’, s_5 , of nearest neighbours most probably represents the true position of the tracer, the algorithm first determines the centroid along l_5 of the 6 relevant points, \vec{P}_n (see Fig. 18).

Considering again our assumptions, one can reasonably assume that all true LoRs should fall within a distance R from the calculated centroid. Any LoRs not satisfying this constraint are then classified as corrupt. For the case of set s_2 , LoRs l_{10} and l_{11} that were previously part of the set are now discarded, leaving 4 LoRs which we may presume to be uncorrupted (see Fig. 18).

This process is repeated for the sets, s_n , belonging to all other LoRs, with the set possessing the largest number of LoRs taken to represent the uncorrupted data. In cases where multiple sets possess the same number of remaining LoRs, the algorithm simply chooses the first such set. It is typically unimportant which is chosen, as all uncorrupted LoRs belonging to the chosen set are also members of the other sets with the same maximum number of remaining LoRs. The tracer position is then determined using exclusively LoRs from the chosen set.

The use of *a priori* information regarding particle size, geometry and tracer distribution is both the greatest benefit and the greatest drawback of the spherical density method. While the use of this information effectively ensures the absence of corrupt LoRs, its requirement means that the method is currently restricted to a very specific range of tracer geometries. Nevertheless, as the majority of PEPT work to date has employed spherical or near-spherical tracers, this is not a severe limitation in most cases.

In order to perform multiple particle tracking, the algorithm is run as described above, and the centroid of the set containing the largest number of uncorrupted LoRs

CONTENTS




Figure 18. *Left:* The same hypothetical set of LoRs depicted in Fig. 17, showing also a green circle of diameter $2R$ placed at the centroid of all points of closest approach lying along line l_5 (see main text). *Right:* The table from Fig. 17 updated to include only cases satisfying the additional constraint that all LoRs counted within the set s_i must pass within a distance R of the centroid position of l_i .

is taken as the position of the first tracer. The LoRs corresponding to this set are then removed from the data set, and the algorithm applied again to the remaining LoRs. This process is then repeated until all N_t tracers within the system have been detected successfully.

To associate computed centroids from a given sample A with a subsequent sample B , the algorithm computes the absolute separation distances between each new centroid and those from the previous batch. If the timestep between successive frames is suitably small, and the particle velocities suitably low, it can be assumed that for two successive computed centroids, the particle would not have moved significantly from its initial computed position. On this basis, the algorithm can associate each newly computed centroid with the closest neighbour from the previous set. This process is repeated to match all centroids from sample B to a correspondent centroid from A .

An interactive Python notebook illustrating the key steps in the Spherical Density Method algorithm may be found [here](#).

Strengths	Weaknesses
Extremely effective in identifying corrupt LoRs	Requires <i>a priori</i> knowledge of tracer geometry and positron emitter distribution
Mesh-free method	Requires <i>a priori</i> knowledge of N_t
Relatively computationally efficient	Currently only implemented for spherical tracers

Table 6. Major strengths and weaknesses of the spherical density method.

CONTENTS

33

5.8. PEPT using Machine Learning (PEPT-ML)

The PEPT-ML algorithm [17, 18], developed by Nicușan and Windows-Yule at the University of Birmingham uses, as its name implies, machine-learning-based clustering techniques to locate and track PEPT tracers. The PEPT-ML algorithm begins by subdividing the acquired set of LoRs into samples of equal length - ideally about 200 LoRs per sample, but many fewer may be used to maximise temporal resolution; indeed, PEPT-ML has been shown to successfully locate particles using fewer than 20 distinct LoRs. For each sample, all pairs of LoRs whose separation is smaller than a user-defined value d_{max} are transformed into ‘cutpoints’; a cutpoint is simply the half-way point of the segment connecting a pair of LoRs at their point of closest approach. As depicted in Figure 19, cutpoints are denser around tracers. The Hierarchical Density Based Spatial Clustering of Applications with Noise (HDBSCAN) [115] algorithm (an overview of which is given in Appendix 7) is used to extract clusters from a given sample of cutpoints. Unlike previously-described algorithms, where the weighted centre of each cluster is then taken as an approximate tracer position, PEPT-ML considers instead ‘cluster exemplars’, or the most significant intra-cluster data points, to compute the final tracers’ locations. As these exemplars are uniformly distributed along a cluster of cutpoints, this allows PEPT-ML to avoid commonly-encountered ‘parallax’ effects, where the asymmetric LoR (and thus cutpoint) distribution for tracers positioned away from the scanner centre creates clusters with higher densities of cutpoints toward the scanner centre, thus skewing the conventionally-calculated centre. An example of two tracers’ tracked positions after this initial (‘first-pass’) clustering is depicted in Figure 19(d).

Following this initial clustering, the cluster centres found are then themselves clustered *again*, yielding tight, highly accurate tracer trajectories (*two-pass* clustering – see Figure 19 (e)). By windowing the raw PEPT data to create large numbers of ‘overlapping’ first-pass clusters, two-pass clustering is able to retain even sharp, sudden changes in the particle movement.

As the number of cutpoints is quadratically proportional to the local LoR density, the number of cutpoints belonging to a given cluster can be taken as an indirect measure of tracer activity; PEPT-ML uses this ‘tracer signature’ to aid the reconstruction of tracer trajectories, even in cases where particles collide – a potentially powerful ability. The process used by PEPT-ML to link and distinguish particle trajectories involves two stages: a spatio-temporal ‘segregation’ step and a signature-based *linking* step. The segregation stage involves a single-linkage clustering algorithm which constructs a minimum spanning tree (MST, or minimum distance path) connecting all points found, then simply “cuts” all paths longer than a maximum distance. However, points that are in the same spatial position at different times must be differentiated - a custom spatio-temporal metric is therefore used, such that only locations close in the time-sorted array of positions can be connected in the MST. Figure 19(f) depicts the segregated positions found after the second pass of clustering, colour-coded by the trajectory ID assigned.

CONTENTS

34




Figure 19. Illustration of steps used by PEPT-ML in tracking two 1 mm diameter MCC tracers in a bubbling fluidised bed, imaged at the University of Birmingham. Panel *a*) depicts 400 LoRs from the data set; panel *b*) shows their corresponding cutpoints, computed for pairs of LoRs which are closer than $d_{max} = 0.5$ mm. Panel *c*) colour-codes them by the cluster IDs assigned by HDBSCAN: green represents noise, purple is cluster one and orange is cluster two; the cluster exemplars' (see main text) centroids are shown as large circles, blue and red, respectively. Panel *d*) depicts the tracer locations found after the first pass of clustering of 80,000 LoRs. Panel *e*) shows the corresponding tracer locations after the second pass of clustering; both panels *d*) and *e*) colour-code points by their cluster size. Panel *f*) illustrates the separated tracer trajectories, colour-coded by the points' trajectory ID.

If tracers collide or intersect, the signatures of the interacting particle can be used to correctly link the pre- and post-collisional 'ends' of each trajectory.

An interactive Python notebook illustrating the key steps in the PEPT-ML algorithm may be found [here](#). A complete implementation of the PEPT-ML method may also be found on Positron Imaging Centre github repository, <https://github.com/uob-positron-imaging-centre/>.

5.9. PEPT using Expectation Maximisation (PEPT-EM)

The PEPT Expectation-Maximisation algorithm [116] was developed at the University of Edinburgh in collaboration with the University of Birmingham. The method expands on the original concept of the Birmingham Method by clustering lines into 'inlier' and 'outlier' sets. Using a maximum-likelihood approach [117, 118], the LoRs in a set of

CONTENTS

Strengths	Weaknesses
Does not require <i>a priori</i> knowledge of number of tracers	Algorithm parameters can be less intuitive
Can track at least 128 tracers	High memory usage for extreme numbers of tracers
Includes tracer signature-based trajectory separation	
High-performance, parallelised implementation with Python interface	

Table 7. Major strengths and weaknesses of PEPT-ML.

data can be allocated to one of many clusters which are iteratively recalculated until convergence, with LoRs that are due to a scattered coincidence or random coincidence being assigned to a larger ‘*outlier*’ cluster approximating the volume of the detector field of view, and ‘true’ LoRs being attributed to a smaller ‘*inlier*’ cluster corresponding to the tracer.

Unlike other methods, PEPT-EM is based on a physical model of positron emission and LoR generation. For a set of K particles, each particle has a centre defined as \vec{x}_k ($k = 1, 2, \dots, K$) and emits positrons that annihilate a short distance away from the centre at y_k , where the probability of annihilation at y_k is defined by a Gaussian distribution around the centre \vec{x}_k with a variance σ_k^2 . The relative weight of each cluster is defined by ρ_k , loosely characterising the relative activity of each particle.

This sets up the problem as a Gaussian mixture [119], which can be solved using expectation maximisation (EM) [120]. We also denote an outlier cluster where $k = 0$, centred at the middle of the field of view and with a variance of the order of the size of field of view. We call the inverse of this variance α , such that $\alpha^{-3/2}$ is of the order of the volume of the field of view. In practice, we use this value as a constant that requires some fine tuning. The weight of this cluster, ρ_0 , is also defined by normalising ρ such that $\sum_{k=0}^K \rho_k = 1$.

Using this information, we define the log-likelihood of finding each cluster with parameters $\vec{x}_k, \sigma_k^2, \rho_k$ in a set of lines \mathcal{L} as

$$L(\{\vec{x}_k, \sigma_k^2, \rho_k\}_{1 \leq k \leq K} | \mathcal{L}) = \sum_{\ell \in \mathcal{L}} \log \left(\rho_0 \alpha + \sum_{k=1}^K \rho_k \sigma_k^{-2} e^{-D^2(\vec{x}_k, \ell)/2\sigma_k^2} \right). \quad (14)$$

The log-likelihood is preferred as it is a concave function and therefore easier to maximise than the likelihood. However, maximising the likelihood explicitly is difficult due to the large number of variables, hence we use expectation-maximisation. With an initial guess of the clusters, we are able to optimise and improve the solution iteratively until we maximise the likelihood.

In order to use expectation-maximisation, we introduce unobserved, or latent, variables. We can think of this variable as imaginary, used to simplify our model.

CONTENTS

36

In this case, the latent variable we use is ζ and is associated with each LoR and labels the LoR as being associated with one of K clusters, such that if the LoR was generated by a positron annihilation originating from the particle k , then the value of $\zeta_\ell = k$ where $\zeta_\ell \in 0, 1, \dots, K$.

We can then, for each line and each cluster k , determine the probability that $\zeta_\ell = k$, given as the latent weights.

$$w_{\ell,k} \left(\left\{ \vec{x}_{k'}, \sigma_{k'}^2, \rho_{k'} \right\}_{1 \leq k' \leq K} \right) = \frac{\rho_k \sigma_k^{-2} e^{-D^2(\vec{x}_k, \ell)/2\sigma_k^2}}{\rho_0 \alpha + \sum_{k'=1}^K \rho_{k'} \sigma_{k'}^{-2} e^{-D^2(\vec{x}_{k'}, \ell)/2\sigma_{k'}^2}} \quad (15)$$

for $1 \leq k \leq K$, where $k = 0$ is the outlier cluster and $w_{\ell,0} = 1 - \sum_{k=1}^K w_{\ell,k}$ and $D^2(\vec{x}, \ell)$ is the minimum distance between the cluster centre \vec{x} and the line ℓ , where \vec{y} is a point on the line:

$$D^2(\vec{x}, \ell) = |\vec{x} - \vec{y}|^2 - ((\vec{x} - \vec{y}) \cdot \vec{n})^2, \quad (16)$$

Note that calculating the latent weights requires an initial guess for the centroids \vec{x}_k , variance σ_k and weights ρ_k of our K clusters. Once we have the latent weights calculated, we use them to recalculate \vec{x}_k, σ_k and ρ_k using

$$\vec{x}_k = \mathbf{Centroid}_k \equiv \left(\sum_{\ell \in \mathcal{L}} w_{\ell,k} \mathbb{T}_\ell \right)^{-1} \sum_{\ell \in \mathcal{L}} w_{\ell,k} \mathbb{T}_\ell \vec{y}_\ell, \quad \text{with } \mathbb{T}_\ell = \mathbb{I}_3 - \vec{n}_\ell \otimes \vec{n}_\ell. \quad (17)$$

where \mathbb{I}_3 is the identity matrix and \vec{y}_ℓ and \vec{n}_ℓ are a point on the line ℓ and the unit vector along ℓ respectively. The optimal variances σ_k^2 are given by

$$\sigma_k^2 \equiv \text{Var}_k \equiv \frac{\sum_{\ell \in \mathcal{L}} w_{\ell,k} D^2(\mathbf{Centroid}_k, \ell)}{2 \sum_{\ell \in \mathcal{L}} w_{\ell,k}}. \quad (18)$$

Finally, the optimal component weights ρ_k are given by

$$\rho_k = R_k \equiv \frac{1}{N} \sum_{\ell \in \mathcal{L}} w_{\ell,k}, \quad (19)$$

with N the number of LoRs. The weight of the outliers is determined as $\rho_0 = 1 - \sum_{k=1}^K R_k$.

The iterative nature of the algorithm requires alternating between the calculation of the latent weights, and maximising the likelihoods using equations 17, 18 and 19 until convergence is achieved, as shown in Algorithm 1. The convergence of an inlier and the outlier cluster is illustrated in figure 20.

The number of iterations to reach convergence can vary depending on the number of clusters and the initial guess. Typically, for the first frame, we apply of the order of 1000 iterations. However, when tracking trajectories, we can achieve convergence using fewer iterations provided that the time interval between batches of LoRs is sufficiently

CONTENTS

37




Figure 20. *Left:* A schematic of the PEPT-EM model of an LoR ℓ being emitted at point y around the surface of a Gaussian cluster centred on x with a variance σ^2 . *Right:* A simple schematic showing the evolution of the clusters generated by the PEPT-EM algorithm, showing clusters $k = 0, 1, 2, 3$. The large outlier cluster, $k = 0$, accounts for detected LoRs that have undergone scatter or are due to random coincidences, while smaller clusters converge using their most-likely LoRs.

Algorithm 1 PEPT-EM

- 1: Initialize $\{\mathbf{x}_k, \sigma_k^2, \rho_k\}_{1 \leq k \leq K}$
 - 2: $\rho_0 \leftarrow 1 - \sum_{k=1}^K \rho_k$
 - 3: **repeat** ▷ Main loop
 - 4: compute $w_{k,\ell}$ from (15) ▷ Expectation step
 - 5: **for** $1 \leq k \leq K$ **do** ▷ Maximization step
 - 6: $\mathbf{x}_k, \sigma_k^2, \rho_k \leftarrow \text{Centroid}_k, \text{Var}_k, R_k$
 - 7: $\rho_0 \leftarrow 1 - \sum_{k=1}^K \rho_k$ Me
 - 8: **end for**
 - 9: **until** convergence
-

small that the particle centroids have only moved a small distance. Consequently, by updating the position of individual clusters we implicitly achieve trajectory separation.

A further feature of PEPT-EM, which isn't used in this paper, is that one may take a Taylor expansion of the particle centroids within a single batch of LoRs \mathcal{L} with respect to time, to account for the motion and acceleration of particles within a single batch of LoRs. As such, not only can we maximise the likelihood of each cluster k having a centroid \vec{x}_k , but also a velocity $\delta \vec{x}_k / \delta t_{\mathcal{L}}$ and acceleration $\delta^2 \vec{x}_k / \delta^2 t_{\mathcal{L}}$. This can significantly reduce the amount of post-processing needed on PEPT data in order to produce velocity vector fields, for example, and can significantly assist in improving the accuracy of tracking.

An interactive Python notebook illustrating the key steps in the PEPT-EM algorithm may be found [here](#).

6. PEPT Methodologies – a Quantitative Comparison

As is apparent from the preceding section, there exist many, strongly divergent PEPT methodologies. Although all documented in peer-reviewed literature, it is nonetheless

Strengths	Weaknesses
Can track a large number of tracers Relatively computationally efficient Includes simple trajectory separation Extended to track velocity & acceleration	Number of iterations to convergence can be arbitrary Requires <i>a priori</i> estimate of N_t Tuning of parameter α is tricky With strongly differing activity, multiple clusters may converge on same point

Table 8. Major strengths and weaknesses of PEPT-EM.

impossible to compare their effectiveness directly and reliably from this literature alone, due largely to two reasons: firstly, the different algorithms have been developed and tested using different detectors and tracers, meaning quoted values of spatial and temporal resolution quoted are system-specific. Secondly, the tests used to benchmark said algorithms vary widely between different publications.

A major aim of this work is to provide a comprehensive and rigorous standardised testing framework for the benchmarking of PEPT algorithms. In the subsections below, introduce the proposed framework, and apply it to the methodologies described in Section 5, focusing specifically on the most modern and/or widely-used algorithms.

6.1. A Standard Testing Framework for PEPT

The framework has been carefully designed to allow a comprehensive, quantitative and, importantly, fair comparison of diverse PEPT algorithms. Firstly, as different algorithms will inherently possess different strengths and weaknesses, the benchmarking framework comprises a number of individual tests, each designed to evaluate a specific aspect of an algorithm’s performance. As not all PEPT algorithms are designed for multiple-particle tracking, it is divided into two parts, one assessing an algorithm’s ability to locate a single tracer, the other its ability to locate – and, crucially, distinguish – multiple particles. Secondly, our analysis judges an algorithm’s ability based on a number of distinct quantities, including both spatial and temporal resolution, the number of particles it can successfully track, its resilience to noise, and various other factors. Thirdly, each individual test must be performed using two distinct camera geometries (based, respectively, on the ADAC Forté [14, 21] and ECAT EXACT [121, 122]), such that algorithms developed for a ring camera are not unfairly disadvantaged by being tested exclusively on a planar camera, and vice-versa[‡]. Finally, all tests are produced using GATE simulations (as described in Section 4) such that the ‘true’ particle positions

[‡] As the scanner geometries simulated are based on real cameras currently used in PEPT imaging, the geometries, sensitivities, and other key detector properties will inherently differ between the two – i.e. the absolute accuracy achievable on the ADAC will differ from that achievable on the ECAT. As such, the results of different algorithms should only be compared in absolute terms for each individual scanner separately.

CONTENTS

39

against which measured locations are compared can be known to an arbitrary degree of accuracy and precision.

It is important to note that the aim of this comparison is not to present a single ‘best’ algorithm, but rather to highlight the relative strengths and weaknesses of each methodology, such that future researchers may choose the most suitable method – or indeed methods – for their specific goals. To ensure a lack of bias, the tests for each algorithm were performed by the original developers thereof, and none of the participants were permitted to view the ground-truth data until after the submission of their results.

The benchmarking tests, alongside all data presented herein, have been made available in an open-source repository [here](#), allowing the developers of future PEPT algorithms to assess their capabilities against contemporary methods easily and meaningfully, thus providing a much-needed standard for the PEPT technique.

For all simulated tests, unless otherwise stated, 10 MBq of fluorine-18 is prescribed as the source activity for a 2 mm diameter, volumetrically activated glass tracer. The activity is defined as positron-emission.

6.2. Single Particle Tracking Tests

In this first suite of tests, we evaluate the fundamental abilities of a PEPT algorithm: its achievable spatial and temporal accuracy, its ability to track rapidly-moving tracers, its ability to mitigate noise, and its consistency throughout a system’s field of view. We include in this test the two algorithms which, historically, have received the most use (the Birmingham method and the G-means method) as well as the most recent algorithms from each of the world’s main PEPT centres (SBSR, FPI, the Spherical Density method, PEPT-ML and PEPT-EM).

6.2.1. Minimum activity test– The first, and most fundamental, test is designed to evaluate the maximal accuracy of an algorithm, as well as its functionality using small numbers of LoRs (i.e. locating low-activity tracers). The test simply requires the reconstruction of a single tracer, using a number of LoRs $N_s \in [10^1, 10^4]$. For each sample size, N_s , 10 repeat tests are performed, each with a different random set of LoRs.

The performance of a given algorithm is judged on the following criteria: i) the minimal static location error, σ_s , achieved for any N_s ; ii) the smallest number of LoRs for which a particle can be located successfully. For the case in which, in test ii), multiple algorithms share the same minimum sample size, $N_s^{min.}$, they are differentiated based on the location error achieved at $N_s^{min.}$. For a given N_s , the location error is calculated as the mean over all N_{rep} repeat tests of the absolute distance between the true particle location, \vec{r}_t , and the measured particle location, \vec{r}_m , i.e.

$$\sigma_s = \frac{1}{N_{rep}} \sum_i^{N_{rep}} |\vec{r}_t - \vec{r}_m| \quad (20)$$

CONTENTS

40




Figure 21. The error, as calculated from equation (20), in PEPT detected positions decreases when more LoRs are used. Data is shown for both the ADAC forte detector geometry (A) and the ECAT EXACT detector geometry (B). The error bars shown represent the standard deviation of the results achieved across multiple repeat tests.

As one may intuitively expect, location error increases monotonically with decreasing sample size for all algorithms tested. Within their respective margins of error, the precision of all algorithms shows a similar power law relation with N_s . Despite being the oldest algorithm, the Birmingham method in fact achieves the highest overall accuracy in both camera geometries, and performs consistently well across all tests. In many – indeed a majority of – cases it outperforms even the most recent algorithms, though PEPT-ML and the FPI method yield superior results in some cases.

That the Birmingham method performs so well in this test – in particular for high numbers of LoRs – is perhaps not surprising, as its iterative nature means that, for static tracers such as those imaged here, it can hypothetically achieve arbitrarily high accuracy [14]).

In terms of the various algorithms' abilities to discern the position of low-activity tracers (point ii) above), 5 of the 7 algorithms tested managed to reliably provide a location using only 10 LoRs, though in all cases the location error was > 1 mm. For $N_s \leq 50$, the Spherical Density Method did not return a result; while the G-means algorithm returned values, these were in many cases false locations, hence the significant error bars shown. All algorithms were found to be able to yield sub-millimetre accuracy with as few as 100 LoRs.

6.2.2. Maximum velocity test– This test is intended to establish the ability of a PEPT algorithm to track fast-moving particles. The test involves reconstructing the trajectory of a moving tracer following a one-dimensional sawtooth path (Fig. 22). The tracer velocity varies from 0.04 m/s to 20 m/s, a span designed to represent, and indeed exceed, the range of velocities one may reasonably expect to encounter in a real PEPT

CONTENTS

41

experiment. The sawtooth path is chosen to introduce significant accelerations and discontinuities, thus further challenging the algorithms. The tracer activity is fixed at 10 MBq, a value chosen to be a) representative of values used in experiment, b) adequately high that the particle motion can, hypothetically, be tracked in all cases, and c) adequately low that neither positron camera is expected to experience significant dead time [21].




Figure 22. Example trajectory for the velocity tests over two periods of motion. A: Forte geometry and for $v_{tracer} = 80\text{mms}^{-1}$ B: Forte geometry for $v_{tracer} = 10240\text{mms}^{-1}$. C: ECAT geometry and for $v_{tracer} = 80\text{mms}^{-1}$ D: ECAT geometry for $v_{tracer} = 10240\text{mms}^{-1}$.

The phenomenology observed here is somewhat richer than in the preceding section. The Birmingham algorithm, though typically producing the best results in the previous, stationary-source test now exhibits the largest location errors in a significant majority

CONTENTS

42

of cases. Conversely PEPT-EM, which produced some of the weakest results for the stationary source in the ECAT geometry now produces the smallest errors for all but two cases in the same geometry. For the ADAC, however, while PEPT-EM again performs strongly for slow-moving tracers its accuracy falls off rapidly as velocity increases, to the point where at high velocities it gives the largest error of all algorithms tested. This is probably due to the fact that at the highest velocities, where only relatively few LoRs per point can be suitably used, the algorithm cannot meaningfully identify both an ‘outlier’ and ‘inlier’ cluster. PEPT-ML, meanwhile, consistently performs either best or second-best across both camera geometries and all velocities tested. The SBSR algorithm also performs extremely well, in particular in the ADAC geometry, and for the case of fast-moving tracers – as may perhaps be expected from its intended purpose of accurately tracking particles using minimal numbers of LoRs.

The majority of algorithms approximately obey a power-law relation between tracer velocity and location error. The increase in tracer speed is expected to cause a greater spread of LoRs within each reconstruction timestep and thus a decrease in reconstruction accuracy. In response to increasing tracer speed, it is typically helpful to decrease the reconstruction timestep to mitigate this effect. However, this inherently leads to a reduced total number of LoRs available for each tracer location, thus reducing spatial resolution. Users in this exercise were left to balance these competing defects in the manner they deemed most suitable for their own respective algorithms. In particular, the decrease in location accuracy with a decreasing numbers of LoRs per location is clearly observed in the results for the G-means and FPI algorithms: the results for these algorithms show apparent step changes in error, occurring approximately every other data point. The fact that these step-changes align with changes in the reconstruction timestep implemented by the user strongly suggests that it is indeed the changing number of LoRs used for reconstruction causing this effect. A similar effect may be observed in the PEPT-EM results, with changes in accuracy at each datapoint.

Considering the above-discussed trade-off between temporal and spatial resolution, it is potentially useful to reframe our data in a way which compensates for this effect. In Fig. 24, rather than the pure spatial resolution achieved by the algorithms – which takes no account of *how often* a particle was detected – we consider instead the ‘*spatiotemporal resolution*’, defined as the average location error divided by the root of the detection rate^{††}. In terms of this metric, the PEPT-EM algorithm – which uses comparatively large amounts of data to produce particularly high location accuracy – performs less well, while the abilities of PEPT-ML and the SBSR methods – both designed explicitly to handle low-activity tracers – are more strongly emphasised; possible reasons underlying this, omitted here for brevity, are discussed in detail in Appendix 8. For both the spatial and spatiotemporal cases, the SBSR algorithm is observed to perform particularly well for high tracer velocities.

^{††}The square root is chosen as the most reliable normalising factor as it is well known that the precision of PEPT location varies approximately as $\frac{1}{\sqrt{N}}$ where N is the number of LoRs used per particle location event [14].

CONTENTS




Figure 23. The error in PEPT detected positions increases when tracers move at a higher velocity. A: Forte detector geometry and movement in X-direction. B: Forte detector geometry and movement in Z-direction. C: ECAT detector geometry and movement in X-direction. D: ECAT detector geometry and movement in Z-direction.

6.2.3. Scatter sphere test— The goal of this test is to establish the ability of the different algorithms tested to handle noisy data. In order to produce noise, the GATE-simulated source is placed at the centre of a 50 mm diameter sphere whose physical properties are varied to elicit different levels of both scatter and attenuation, each of which may act to increase the presence of ‘random’ photons and thus noise in a given PEPT measurement. The scatter media used are summarised in Table 9.

The materials used are chosen to provide a significant range of attenuation coefficients and thus degrees of both attenuation and scatter, from air – in which almost zero scatter is expected – to BGO, in which only one in every thousand gamma ray pairs can be expected to reach the detectors unimpeded. The spheres are offset vertically by 20 mm from the centres of both cameras so as to avoid any non-representative data

CONTENTS

44




Figure 24. The calculated spatiotemporal resolutions of PEPT detected positions increases when tracers move at a higher velocity. A: Forte detector geometry and movement in X-direction. B: Forte detector geometry and movement in Z-direction. C: ECAT detector geometry and movement in X-direction. D: ECAT detector geometry and movement in Z-direction.

arising due to the special case of a perfectly-centred source.

Unlike the two preceding tests which produced relatively clear trends, those yielded by the scatter sphere test are more complex. While the overall trend across all data sets is, unsurprisingly, a general decrease in accuracy with increasing density (and thus increasing scatter/attenuation), the individual behaviours of several algorithms are observed to be somewhat erratic.

Only the G-means algorithm yields a continuous monotonic trend, suggesting it is the least sensitive to noise; the downside, however, is that on average it yields the lowest accuracy. Though non-monotonic, the FPI, Spherical Density and SBSR methods all yield relatively consistent and strong results across both detector geometries. The

CONTENTS




Figure 25. Left: Image of the 100 mm diameter scatter sphere and three gamma ray pairs, one of which has scattered with the sphere material. Centre: Image of the scatter sphere inside the ADAC Forte detector. Right: Image of the scatter sphere inside the ECAT geometry.

Material	Density (g/cm ³)	Cross-section (cm ²)	Attenuation Coeff. (cm ⁻¹)
Air	0.00129	0.0806	0.000111
Polyethylene	0.960	0.0986	0.0946
Water	1.00	0.0960	0.0960
PVC	1.65	0.0890	0.147
Aluminium	2.70	0.0837	0.226
NaI	3.67	0.0933	0.342
CZT	5.68	0.0896	0.509
Steel (SS304)	7.92	0.0832	0.659
Copper	8.96	0.0827	0.741
BGO	7.13	0.1501	1.070

Table 9. List of materials and properties used in the scatter sphere tests.

Birmingham method shows both relatively good consistency and accuracy (indeed, the best accuracy for the highest attenuation) in the ADAC geometry, but is somewhat less reliable in the ECAT.

These findings raise an interesting question as to whether ability or reliability is the most important factor when choosing an algorithm for an experiment involving significant variability in the degrees of scatter and attenuation likely to be encountered. If the former, PEPT-EM would be the most sensible choice, but would run the risk of producing extreme outliers. If the latter, then the FPI, Spherical Density or SBSR method would provide a safer choice.

The extreme variation observed is most likely explained by the fact that under high-scatter conditions, there exists an increased probability that random fluctuations in the LoR density field may be comparable in magnitude to those registered at the true particle location, and that certain algorithms are simply more susceptible to these false locations. A deeper investigation into the origins of the observed inconsistencies in location would make a valuable future study.

CONTENTS

46




Figure 26. The error in PEPT detected positions increases when more LoRs are corrupted by scattering. A: Forte detector geometry. B: ECAT detector geometry.

6.2.4. *Field of view test*— In this test, a simple, linear tracer trajectory is simulated, with the particle starting outside the camera at one end of its field of view, and passing through its centre and out the other end. The goal of this test is to establish the ability of the algorithms tested to successfully track particles not only in the centre of a camera's field of view but also in other regions, where sensitivity and efficiency are reduced.




Figure 27. Left: Forte geometry. Right: ECAT geometry. Length of the trajectory for the FOV tests.

As is immediately clear from Fig. 28, for the ADAC geometry, the Stanford SBSR method markedly outperforms other algorithms both in terms of its peak accuracy and consistency across the bulk of the detector face. While all other algorithms experience

CONTENTS

a sharp increase in their error margins as the tracer leaves the most sensitive central region of the camera, an error of < 0.2 mm is achieved across the vast majority of the trajectory for SBSR. Despite this high accuracy, the code also offers very impressive temporal resolution, as may be seen from the density of points in Fig. 28, again in line with its intended purpose of providing high-resolution data using relatively low numbers of LoRs. Interestingly, this excellent performance is not reflected in Fig. 29, as the overall mean error is significantly skewed by the presence of large deviations encountered at the very edges of the scanner (not shown in Fig. 28 so as to maintain a sensible y -axis scale). This is a salient reminder that such simple scalar statistics – while useful for making quick, broad comparisons – can, in some cases, be misleading. Where this is found to be true in the present work, as in this case, we have been careful to include additional data to ensure a fair and meaningful comparison.




Figure 28. As the tracer nears the edges of the FOV, the error increases. A: Forte geometry. B: ECAT geometry. Every fifth detection is plotted and moving average filter of 20 points is applied to smooth the plotted trajectories for each algorithm.

In the ECAT geometry, while at its peak near the centre of the system the SBSR algorithm still achieves the highest accuracy of all codes tested, it is the FPI algorithm that demonstrates the greatest consistency across the full axial length of the detector. The lower accuracy achieved by the SBSR algorithm in this case is most probably due to the relatively short bore of the ECAT camera – as the locations determined by the SBSR algorithm take into consideration both previous and subsequent timesteps, it is possible that high errors achieved near the edge of the system are acting to increase the error also of other points further within the detector field of view that also rely on this information.

In terms of overall accuracy, all codes are observed to achieve sub-millimetre precision for both camera geometries. In the ADAC geometry, the Birmingham and PEPT-ML algorithms are observed to provide the highest accuracy; in the ECAT

CONTENTS

48

geometry the FPI and PEPT-EM algorithms perform similarly well. Where the Birmingham, FPI and PEPT-EM algorithms all follow relatively consistent trends, however, PEPT-ML's accuracy varies much more erratically, though this may be due to the small sample size used, allowing it to yield a significantly higher temporal resolution than the other codes whilst maintaining a comparable spatial resolution – as is clearly illustrated in Fig. 29.

A key observation from this test is that the fundamental tracking philosophy of whether a single frame contains a fixed number of LoRs or occurs over a fixed timestep can considerably impact the results obtained. At the edges of the field of view, the camera sensitivity varies rather significantly. In this case, the Birmingham algorithm, for example, simply discards data points at the edges of the system due to the absence of a suitable number of 'good' LoRs for it to locate a particle at each end of the trajectory – the calculated uncertainty on the minimum distance point in these cases is above the defined threshold, meaning that the points are discarded. However, for the most part of the trajectory the Birmingham algorithm is perhaps over-sampling, with a location every 250 LoRs, meaning that the tracked locations appear relatively noisy.

In contrast, PEPT-EM provides a relatively smooth trajectory towards the centre of the detector, sampling at a fixed rate every 10 ms and thus utilising approximately 8000 LoRs per location; however at each end there is a deviation from a straight line, as fewer LoRs are used at each location. For the ECAT dataset, if the first 10 and final 10 points are excluded, the RMS error to a straight line fit is 0.11 mm for the Birmingham method, and 0.025 mm for Expectation-Maximisation. However, if we include all of the points the RMS error becomes 0.13 mm for the Birmingham method and 0.18 mm for Expectation-Maximisation. If the number of LoRs per location is increased to 5000 for the Birmingham Method, then we reduce the RMS error to 0.025 mm across all points, at the expense of tracking the ends of each trajectory. Therefore, for problems like this in which the tracer explores the entirety of a scanner's field of view, the user should consider whether they value tracking the entirety of a trajectory with slightly less accuracy, or choose to use an algorithm that may provide a shorter section of the trajectory but not compromise on resolution.

The G-means and FPI methods similarly use fixed timesteps for reconstruction, and suffer from a loss of accuracy at the axial edges of each scanner, where reduced sensitivity leads to fewer LoRs per timestep. In the case that these regions of reduced accuracy may adversely impact the outcomes of a given measurement, an experiment can be designed such that any data taken in these regions can be rejected. Such a 'clipping' of the scanner field of view to its more sensitive regions has indeed already been employed in PEPT measurements using the FPI method [46, 94].

In general, error values corresponding to the majority of codes follow an approximate 'U'-shape, with a minimum at the centre of the detector and a maximum near its edge. This is probably due to the higher proportion of random events near the edges of the detectors, where it is more likely that one half of an annihilation-pair of gamma rays will hit the detector and the other 'miss'. The corresponding increase in

CONTENTS




Figure 29. Above: Mean errors of each algorithm plotted with error bars denoting the standard deviation of errors. A: Forte geometry. B: ECAT geometry. Below: The spatiotemporal resolution returned by each algorithm. A: Forte geometry. B: ECAT geometry. Below: The spatiotemporal resolution returned by each algorithm. A: Forte geometry. B: ECAT geometry.

corrupted LoRs will tend to ‘pull’ the measured location toward the centre of the field of view. Differences in sensitivity and geometric efficiency at different positions within the cameras’ fields of view are also liable to have an effect on the data quality produced by some, if not all, algorithms near the edge of the field of view. Additionally, there exists a slight asymmetry to the ‘U’-shape. This is because in the Forte geometry the detectors are fixed to a large aluminium gantry ring about which they are able to rotate. Similarly, in the ECAT geometry, tungsten shielding is present which acts to partially shield the detectors from rays potentially emanating from other parts of a given system lying outside the FOV. In both cases, the asymmetrically-positioned material increases

CONTENTS

50

the proportion of detected LoRs corrupted by scatter and/or attenuation, and thus disproportionately increases the error in PEPT reconstructed positions for some regions of the field of view as compared to others.

The main observation which can be drawn from this section is the complexity of the value judgement one must undertake when choosing the most suitable PEPT algorithm for a given situation. Here, the algorithm achieving the lowest overall error is not necessarily the same as that which will produce the highest spatial resolution across the majority of the scanner's field of view, nor are either of these guaranteed to yield optimal temporal resolution. As such, the optimal choice depends sensitively on the goals of the experiment to be performed, the detector geometry used, and the size of the system used. This observation presents a particularly emphatic example of a more general finding from the present work – that there really does not exist a single ‘best’ algorithm overall; rather, different algorithms will excel under different circumstances.

6.3. Multiple Particle Tracking Tests

In this second set of tests, we assess our algorithms' abilities to detect, differentiate and track multiple particles, exploring their accuracy, their ability to distinguish nearby tracers, their ability to simultaneously locate large numbers of tracers, and their resilience against false location events. In these tests, only 5 of the 7 algorithms tested in the previous section are included – the SBSR algorithm has yet to be extended to the case of multiple-particle tracking and the Birmingham method, though capable of locating multiple tracers, does not have an in-built functionality for trajectory separation in the case of equally-active tracers.

6.3.1. Minimum separation test – In this test, two 2 mm diameter tracers, each of activity 10 MBq, are separated by a variable distance x . The value of x is varied from 30 mm down to 2 mm (i.e. the point at which the two simulated tracers are in contact). The purpose of this test is to assess the ability of a given algorithm to successfully resolve two particles in close proximity as opposed to falsely merging them into a single location – an explicitly-noted problem with some older algorithms [84].

It is immediately notable that PEPT-EM, PEPT-ML, FPI and the Spherical Density Method are all capable of successfully resolving the two separate tracers up to and including the point of contact. This is a potentially highly consequential finding, as it had previously been assumed that the PEPT technique was incapable of imaging particle collisions. This observation, therefore, potentially opens up a number of valuable new applications for PEPT. Further testing needs to be performed in order to determine the limitations on tracer size for which particle collisions can be successfully resolved.

All five codes tested achieve sub-millimetre accuracy across the majority of the tested separation values. The accuracy levels achieved by all codes are observed to remain relatively independent of the precise particle separation for comparatively large distances ($\gtrsim 5$ mm), sharply increasing after this point in all cases except for PEPT-

CONTENTS




Figure 30. The measured error in PEPT location for the case of two static tracers separated by various centre-to-centre distances ranging from 2 mm (particles in contact) to 30 mm. A: Forte detector geometry and separation in X-direction. C: ECAT detector geometry and separation in X-direction. D: ECAT detector geometry and separation in Z-direction.

EM, which maintains a remarkable degree of precision (typically $< 100 \mu\text{m}$), even to the point of contact. This impressive consistency is perhaps due to the fact that, unlike other contemporary PEPT methods, the Gaussian mixture model underlying PEPT-EM is based on a more comprehensive mathematical model of the physics underlying PEPT. That is, PEPT-EM seeks to maximise the probability of there being exactly 2 tracer clusters, rather than trying to identify N clusters and their locations. It should also be mentioned, however, that PEPT-EM produced only a single data point from all LoRs in the samples provided, with others attempting to produce multiple locations from the same data, which may have had an influence on the absolute accuracy achieved.

It is also of note that in many of these results there appears to be a gap of roughly 0.1 mm between the accuracy achieved by PEPT-EM and PEPT-ML and that of the G-Means and FPI methods. This may be simply due to a small discretization error caused by the use of voxels in the latter methods. For this exercise, both G-Means and FPI used 1 mm grids; it has been shown in optical peak fitting routines (similar to the methodology employed for FPI localization) that the minimum achievable accuracy

CONTENTS

52

is usually on the order of one-tenth of a pixel [123]. In other words, it is possible that the differences observed between voxelised and mesh-free methods is not an innate difference in the algorithms' abilities, but simply due to the specific choice of voxel size implemented.

6.3.2. False positive test– In this test, the algorithms used must successfully locate an array of 4, 8, and 16 particles arranged, respectively, to form a square, a cube, and a tesseract (Fig. 31), geometries deliberately chosen to include high degrees of symmetry to create 'cross-talk' between LoRs from different particles which may be falsely interpreted as particle positions. This test is therefore designed to assess an algorithm's resilience to false detection events.

All geometries created share a vertex length of 10 mm, meaning that no two particles lie closer than this distance. This specific minimum separation distance is chosen to be greater than the detector resolution and – as demonstrated in the previous section – corresponds to the separation above which the majority of codes' spatial resolutions become invariant with distance, ensuring that the present test does not unfairly 're-examine' abilities tested previously.

Impressively, all codes successfully located all particles in all cases, with no codes producing false locations.

For all algorithms tested, however, we observe a general decrease in accuracy with increasing tracer number. While one may intuitively ascribe this trend of increasing error simply to an increase in particle number, it is notable that a similar trend is not observed in section 6.3.4, where particle positions are entirely random. This suggests that it is indeed the symmetry of the systems modelled here that produces an additional confounding effect. Specifically, we observe that the measured particle locations are biased toward the centre of the constellation of particles (see Fig. 31).

For the lower tracer numbers and lower degrees of symmetry, all algorithms are observed to achieve sub-millimetre accuracy, with PEPT-ML and PEPT-EM both providing resolution on the scale of 100 microns, and the latter achieving accuracy of below 10 microns in the 4 particle case. As in the previous test, we observe a small difference in behaviour between meshed and mesh-free methods, which may be ascribed to the aforementioned pixelation effects. The relatively increased error in the spherical density method most probably arises due to the iterative nature of the algorithm – as data is being discarded for each additional particle detected, the precision of locations will decrease with increasing tracer number, as noted in prior works concerning iterative PEPT algorithms [82, 83].

6.3.3. Orientation tracking test– This test aims to assess not only an algorithm's ability to locate a tracer at a single point in time, but to follow said tracer reliably across multiple timesteps and, importantly, to distinguish successfully the trajectories of multiple particles following intersecting paths. As such, this is a test not only of




Figure 31. *Above:* Figure illustrating the positions of groups of tracers for the false positive test. The three tests correspond to groups of 4, 8 and 16 particles arranged, respectively, in a square, a cube and a tesseract (the four-dimensional analogue to a cube). Note that in reality each group of tracers are simulated separately and centred at 0, 0, 0 mm. *Below:* The positions returned by each PEPT algorithm for the 16 particle test, though some particles are hidden by symmetry in this figure. Some biasing towards the centre is shown, most notably for the FPI and G-Means methods. The SDM algorithm experiences difficulty distinguishing between tracers when they are highly concentrated. A: Forte geometry. B: ECAT geometry.

CONTENTS

54




Figure 32. The error in PEPT detected positions for the groups of particles illustrated in Fig. 31. A: Forte geometry. B: ECAT geometry.

the raw location accuracy of a method, but also the associated trajectory separation algorithms.




Figure 33. This test uses three particles which remain constantly separated from one another by a fixed distance of 34.5mm (panel (a)) whilst being translated across a spherical surface. The paths followed by the three particles are illustrated in panel (b).

We simulate three particles, each separated from the others by a fixed distance, moving continuously across the surface of a sphere of diameter 40 mm at a constant velocity 0.8 m/s (see Fig. 33). The reconstructed trajectories are compared individually against the ground truth data and the mean deviation (equation (20)) across all measured trajectories used as a measure of the algorithm's capability.

All of the tested algorithms were able to distinguish and link the separate paths followed by the three moving particles successfully – that is to say all yield consistent trajectories with consistent IDs, avoiding the common issues of particles ‘switching’ ID at points of intersection or becoming ‘lost’. Once again, all algorithms showed

CONTENTS




Figure 34. Above: Mean error in reconstructed positions of the three tracers (note that all algorithms correctly linked the positions of tracers into unique trajectories). A: Forte geometry. B: ECAT geometry. Below: Spatiotemporal resolution of the trajectory linking test. A: Forte geometry. B: ECAT geometry.

themselves capable of producing sub-millimetre accuracy even for the case of multiple moving tracers. For both detector geometries, PEPT-ML and PEPT-EM produced both the highest and most consistent accuracy.

Interestingly, despite relatively large errors and the largest variability of all codes tested in the ADAC geometry, in the ECAT geometry the G-means algorithm shows a capability similar to that of PEPT-ML and PEPT-EM, and even outperforms its successor, the FPI method. As discussed previously, however, the simple measure of overall error does not account for the frequency of detection. If we consider instead the spatiotemporal resolution (Fig. 34), we see an improved relative performance from the FPI method. Indeed, by this measure, the FPI method also (for the ECAT geometry) outperforms the PEPT-EM method which, despite high overall accuracy, offers a lower temporal resolution.

6.3.4. Large tracer number test– The eighth and final test is intended to assess limitations on the number, N_t , of tracers that can be successfully tracked by a given algorithm, evaluating also how the algorithm’s accuracy is influenced by increasing N_t .

The test consists of 9 distinct data sets, each containing a different number of

CONTENTS

56

tracers placed in random positions throughout the scanners' fields of view. The tests are deliberately conducted with pseudo-random numbers of particles to prevent those attempting the benchmarking tests from making 'educated guesses' as to the expected number of locations.

Two aspects of the codes' performance are directly tested: their ability to locate large numbers of tracers with no *a priori* information regarding the specific number to be located, and the accuracy with which said tracers are located (see Fig. 35). Regarding the former criterion, all algorithms tested perform impressively – indeed, for the ECAT geometry the Spherical Density, FPI and PEPT-ML algorithms successfully locate all particles in all tested cases; the last two also receive a perfect score for the ADAC geometry. The G-means algorithm also performs relatively well, typically losing only a relatively small fraction of tracers in comparatively high- N_t cases. The PEPT-EM algorithm performs well for comparatively low tracer numbers in the ECAT geometry, but loses particles in almost all ADAC test cases, the fraction of lost tracers generally increasing with N_t .

In terms of the accuracy with which the successfully-located tracers are found, the algorithms perform similarly, with PEPT-ML and PEPT-EM generally producing the highest overall accuracy, though it must be remembered of course that PEPT-EM in many cases is detecting a smaller fraction of the total particle number present than the other algorithms, thus potentially skewing this metric.

The consistency in average location accuracy exhibited across the full range of tracer numbers explored is striking: the FPI method, for example, can seemingly detect 79 particles just as precisely as it can detect 6. Indeed, all algorithms tested are found to consistently produce millimetre- or submillimetre-scale accuracy across practically all data sets. This observation bodes very well for future research using large numbers of PEPT tracers, a possibility that to date has yet to be properly explored and exploited.

7. PEPT Applications

Having looked in detail at the algorithms and hardware used to perform PEPT we provide, in this penultimate section, some insight into the ways in which PEPT may be usefully applied in various areas of science, industry, and medicine. The section is split into two parts: in the first, we provide an overview of the major contemporary applications of the PEPT technique, highlighting the fields in which it has been widely and successfully applied, and providing a concise reference guide to some of the important work in these areas. In the second, we discuss some of the new, emerging applications of PEPT, several of which are only possible due to the recent advancements discussed in these pages. In this second subsection we also look at where these and further developments may take us in the near future.

CONTENTS




Figure 35. *Above:* The means error of the PEPT detected positions from the multiple particle tracking test compared to the nearest GATE prescribed particle positions A: Forte geometry. B: ECAT geometry. *Below:* The percentage of particles found by each PEPT algorithm compared to the number prescribed GATE particles. A: Forte geometry. B: ECAT geometry.

7.1. Current Applications

We preface this section by noting that the focus of this Review lies on developments of the PEPT technique itself, as opposed to its applications. There exist already a number of resources focusing specifically on PEPT's applications in various of fields (see, for example, [22, 30, 60, 81, 124, 125]). As such, this section does not aim to provide a detailed analysis of prior works using PEPT, but rather a succinct reference guide introducing a number of useful sources, from which the interested reader may then gain further knowledge.

Early work using PEPT focused on the imaging of industrial process equipment, taking advantage of PEPT's unique ability to image large, opaque, fast-moving systems with high accuracy. Particular attention was paid to batch mixing systems [126–128] used widely in the chemical, food and pharmaceutical sectors. The study of mixing and segregation in industrial processes has remained a major area of investigation using

CONTENTS

58

PEPT since its inception [124, 129–131]; more recently, PEPT has additionally been used to study the fundamental dynamics of mixing and segregation phenomena from a pure-physics perspective [3, 132, 133]. PEPT is also used extensively in the study of numerous other key industrial processes including extrusion [134, 135], fluidisation [7, 35–37, 136, 137], granulation [31, 138, 139], die-filling [140, 141], (hydro)cyclone separation [142–144], anaerobic digestion [10, 145], stirred tanks, and milling [8, 9, 69, 146].

In addition to the more applied research described above, in the later 90s, PEPT became more widely adopted in the study of the fundamental physics of particulate and multiphase systems. Due to its improved data acquisition rate compared to other three-dimensional imaging methods [58] such as x-ray CT [147], RIMS [148] or PET [149], and its improved spatial resolution as compared to radioactive particle tracking [150] or magnetic particle tracking [151], PEPT has proved a highly popular means of studying the dynamics of granular gases [152], whose rapid dynamics researchers could not previously capture precisely using conventional techniques. Granular gases, and more generally ‘vibrofluidised’ granular beds [153], provide a valuable canonical system in which to study the fundamental phenomena and dynamics of particulate media [58], making them a common choice for such pure-physics studies. PEPT has been used to study numerous aspects of these systems, including their granular temperature distributions [154] (and the striking non-equipartition thereof [2, 155]), their diffusive behaviours [32, 156], their velocity distributions [33, 157], and phenomena such as jamming [158], segregation [3], and granular convection [1, 4]. PEPT has also been used for research into the fundamental physics of particulate, multiphase and indeed pure-fluid systems in a number of other geometries, most notably fluidised beds [35, 36, 136, 159], rotating drums [133, 160–163] and stirred tanks [5, 164–167].

PEPT’s ability to provide detailed, three-dimensional data make it an extremely valuable and well-established [168] tool for the validation of theoretical and numerical models, being used to complement and validate a variety of discrete element method (DEM) models of systems, from fundamental systems such as vibrofluidised beds [1–4] and rotating drums [133, 169] to comparatively complex industrial equipment including mixers [130, 170–173], mills [8, 146, 174, 175], granulators [176] and other process equipment [177]. PEPT and DEM represent in many ways an ideal complementary pairing: the former provides detailed and accurate three-dimensional information regarding the dynamics of a system, but offers only limited information regarding particle contacts and stresses, and is comparatively expensive to perform, limiting the feasibility of large-scale parameter studies. DEM, meanwhile, provides full information regarding particle contacts and stresses, and can be performed cheaply using a standard PC, making large parameter studies, and the exploration of extreme conditions (e.g. frictionless particles, zero gravity) easily achievable. However, DEM results in isolation, due to the many simplifying assumptions of the technique [58], cannot be trusted to be accurate – or even physical [74]. Rigorous, multi-faceted validation, such as that which may be performed through comparison with PEPT data, is required if their results are to be useful.

CONTENTS

59

The above holds true also for computational fluid dynamics (CFD) [178] and indeed coupled CFD-DEM models [179], widely used, for example, to simulate the dynamics of fluidised beds [180–182]. In the literature one may find many examples of combined PEPT-CFD (and PEPT-CFD-DEM) studies of these systems [183] as well many others, including wet mills [8, 146], hydrocyclones [142, 184], and various forms of agitated vessels, from dishwashers [185] to anaerobic digesters [186], and spanning systems comprising simple Newtonian fluids, non-Newtonian fluids and fluid-solid suspensions [187, 188].

7.2. Nascent and Future Applications

As is clear from the preceding section PEPT has, historically, been used predominantly by physicists, engineers, and in industry. One field in which PEPT has been under-exploited, however, is that of geophysics/geomechanics. While the potential value of the technique in this field has been suggested in prior publications [189–191], and systems widely studied using PEPT, such as rotating drums and vibrated beds, are commonly used as laboratory-scale proxies for geophysical systems [192–194], the *direct* application of PEPT to geophysical problems remains, to date, very limited. Much like in the industrial sector – where it enjoys great popularity – PEPT’s ability to image large, dense, opaque systems with unusual geometries would prove highly valuable in the field of geophysics. Based on this suitability, and the fact that, over the past decade, positron emission *tomography* (PET), has begun to be adopted by the field [195–197], we are hopeful that PEPT may soon follow suit.

In recent years, researchers have also begun to explore the potential of PEPT in the biomedical field [11, 12]. Advances in the ability to label small particles (section 3.2), combined with improved hardware (section 3.1) and the development of new, more efficient algorithms (section 5) mean it is now possible to track the motion of cells and other small objects *in vivo* [95]. Such an ability carries myriad clinical and non-clinical possibilities, for example in emission-guided radiation therapy for the treatment of cancers [198, 199], to track cells during cell-based therapies [200, 201], or in the study of blood flow, for example helping to validate and improve models of flow within ventricular assist devices [202, 203]. It may also potentially be used in conjunction with 4D PET/CT imaging, e.g. to improve image quality by estimating motion due to breathing [204]. Of course these represent just a few examples; there exist many more possible applications of PEPT in the biomedical discipline, making it an exciting field for future exploration.

8. Summary and Future Outlook

It is clear both from the review of recent literature presented and the new work performed in this article that the past decade has seen a significant expansion in both the applications of and capabilities of the PEPT technique. Perhaps our most notable finding is that of the diverse new methods developed, there does not exist a single

CONTENTS

60

‘best’ algorithm; rather, all exhibit particular strengths and weaknesses. As such, to get the most out of the PEPT technique, it is advisable to use different algorithms for different goals and/or in different systems. For example, for the tracking of slow-moving, relatively active single tracers one can make very effective use of the original, highly computationally-efficient Birmingham method; conversely, for applications requiring the tracking of large numbers of particles, or where the imaging of particle-particle collisions is necessary, the more modern, but more computationally intensive, algorithms such as PEPT-ML, PEPT-EM or FPI are of value; if tracer activity is limited meanwhile – e.g. when tracking small, fast-moving particles – the Stanford SBSR method may be the optimal choice. We hope the detailed benchmarks and algorithm overviews provided in this work will be of value to future researchers in choosing the most suitable methodology/ies for their analysis.

Our results have also demonstrated that several contemporary algorithms possess impressive capabilities in the tracking of large numbers of individual tracers at relatively low activities. To date, the vast majority of PEPT research has focused on single particles, with a handful of studies employing small but non-unity numbers of tracers. Combined with our findings that several of these algorithms can also detect and distinguish contacting tracers, this opens up a number of exciting possibilities for entirely novel future studies involving large numbers of interacting tracers. These multiple-particle studies will not only lead to more efficient collection of single-tracer data for the calculation of previously attainable flow and transport phenomena, but will also enable the collection of three-dimensional, multiple-point, multiple-time statistics, helping us to more deeply understand the flow behaviours of a diverse range of system. Such data may be otherwise unobtainable in a number of complex and opaque geometries.

Beyond the development of PEPT *algorithms*, the future potentially also holds a variety of exciting developments concerning PEPT tracers and detectors. Regarding the former, a major goal in the field remains the creation of ever smaller tracers, so as to realise the many potentially groundbreaking medical applications discussed in the preceding section, as well as opening new potential avenues of investigation in industrial applications. For example, particles in the 1-20 μm range have increasing industrial relevance in processes such as fluidisation, mineral extraction/processing and mining, recovery of low grade ores, or in various closed cycle processes relating to the green economy, where particle ‘fines’ have increasing relevance to high-intensity or low-efficiency processes. The advent of PEPT algorithms capable of successfully tracking smaller and more multitudinous tracers also gives rise to the possibility of directly studying processes such as grinding and milling in which the tracer particles used actually undergo fracture. Researchers at the University of Birmingham, for example, are currently using PEPT to study the production of graphene via shear-exfoliation [205, 206] using carbon flakes labelled with positron-emitting ^{11}C . Similar techniques, potentially in tandem with the application of Positron Emission Projection Imaging (PEPI) [207], may also be employed to study other industrially-important processes such as dissolution and dispersion.

CONTENTS

61

Regarding detectors, in addition to the general progression in the data acquisition rates of positron cameras discussed in section 3.1, some modern PET scanners can also offer time-of-flight data [208] which, compared to the simple LoRs offered by conventional scanners, provide additional localisation information regarding individual annihilation events. The availability of this additional information has already facilitated a step-change in PET imaging capabilities, and stands to prove similarly transformative in the field of PEPT, where it may be used, for example, to ‘pre-filter’ the raw data fed to the algorithms described in section 5, or to develop entirely new algorithms. PEPT practitioners are also yet to take full advantage of full-body positron cameras [209], which not only offer a larger field of view than conventional ring scanners, but also minimise issues relating to the non-uniform sensitivity of narrower scanners, the negative effects of which have been clearly demonstrated in section 6.

References

- [1] RD Wildman, JM Huntley, and DJ Parker. Convection in highly fluidized three-dimensional granular beds. *Physical review letters*, 86(15):3304, 2001.
- [2] RD Wildman and DJ Parker. Coexistence of two granular temperatures in binary vibrofluidized beds. *Physical review letters*, 88(6):064301, 2002.
- [3] CRK Windows-Yule, T Weinhart, DJ Parker, and AR Thornton. Effects of packing density on the segregative behaviors of granular systems. *Physical review letters*, 112(9):098001, 2014.
- [4] CRK Windows-Yule, N Rivas, and DJ Parker. Thermal convection and temperature inhomogeneity in a vibrofluidized granular bed: The influence of sidewall dissipation. *Physical review letters*, 111(3):038001, 2013.
- [5] YS Fangary, M Barigou, JPK Seville, and DJ Parker. Fluid trajectories in a stirred vessel of non-newtonian liquid using positron emission particle tracking. *Chemical engineering science*, 55(24):5969–5979, 2000.
- [6] P Pianko-Oprych, AW Nienow, and M Barigou. Positron emission particle tracking (pept) compared to particle image velocimetry (piv) for studying the flow generated by a pitched-blade turbine in single phase and multi-phase systems. *Chemical Engineering Science*, 64(23):4955–4968, 2009.
- [7] M Van de Velden, J Baeyens, JPK Seville, and X Fan. The solids flow in the riser of a circulating fluidised bed (cfb) viewed by positron emission particle tracking (pept). *Powder Technology*, 183(2):290–296, 2008.
- [8] Chandana T Jayasundara, RY Yang, BY Guo, AB Yu, Indresan Govender, Aubrey Mainza, Andre van der Westhuizen, and Josh Rubenstein. Cfd-dem modelling of particle flow in isamills—comparison between simulations and pept measurements. *Minerals engineering*, 24(3-4):181–187, 2011.
- [9] LS Bbosa, I Govender, AN Mainza, and MS Powell. Power draw estimations in experimental tumbling mills using pept. *Minerals Engineering*, 24(3-4):319–324, 2011.
- [10] Rebecca C Sindall, Davide Dapelo, Tom Leadbeater, and John Bridgeman. Positron emission particle tracking (pept): A novel approach to flow visualisation in lab-scale anaerobic digesters. *Flow Measurement and Instrumentation*, 54:250–264, 2017.
- [11] Keum Sil Lee, Tae Jin Kim, and Guillem Pratx. Single-cell tracking with pet using a novel trajectory reconstruction algorithm. *IEEE transactions on medical imaging*, 34(4):994–1003, 2014.
- [12] Bernhard Schmitzer, Klaus P Schäfers, and Benedikt Wirth. Dynamic cell imaging in pet with optimal transport regularization. *IEEE transactions on medical imaging*, 2019.

CONTENTS

62

- [13] Marco Jonas Mauritz. Analysis of a dynamic cell imaging model in positron emission tomography. 2020.
- [14] DJ Parker, RN Forster, P Fowles, and PS Takhar. Positron emission particle tracking using the new birmingham positron camera. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 477(1-3):540–545, 2002.
- [15] MR Hawkesworth, CR Bemrose, P Fowles, and MA O’Dwyer. Industrial application of positron emission tomography. Tomography and Scatter Imaging, N McCuaig and R Holt (Eds.), IOP Publishing, Bristol, pages 67–79, 1989.
- [16] MR Hawkesworth, DJ Parker, P Fowles, JF Crilly, NL Jefferies, and G Jonkers. Nonmedical applications of a positron camera. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 310(1-2):423–434, 1991.
- [17] AL Nicușan and CRK Windows-Yule. Positron emission particle tracking using machine learning. Review of Scientific Instruments, 91(1):013329, 2020.
- [18] Savannah Mandel. Machine learning is used to conduct positron emission particle tracking, 2020.
- [19] Seth Langford, Cody Wiggins, Daniel Tenpenny, and Arthur Ruggles. Positron emission particle tracking (pept) for fluid flow measurements. Nuclear Engineering and Design, 302:81–89, 2016.
- [20] Matthew Herald, Zachary Bingham, Roque Santos, and Arthur Ruggles. Simulated time-dependent data to estimate uncertainty in fluid flow measurements. Nuclear Engineering and Design, 337:221–227, 2018.
- [21] Matthew Herald, Tzany Wheldon, and Christopher Windows-Yule. Monte carlo model validation of a detector system used for positron emission particle tracking. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 993:165073, 2021.
- [22] CRK Windows-Yule, JPK Seville, A Ingram, and DJ Parker. Positron emission particle tracking of granular flows. Annual Review of Chemical and Biomolecular Engineering, 11, 2020.
- [23] D Strulab, G Santin, D Lazaro, V Breton, and C Morel. Gate (geant4 application for tomographic emission): a pet/spect general-purpose simulation platform. Nuclear Physics B-Proceedings Supplements, 125:75–79, 2003.
- [24] G Santin, D Strul, Delphine Lazaro, L Simon, M Krieguer, M Vieira Martins, Vincent Breton, and C Morel. Gate: A geant4-based simulation platform for pet and spect integrating movement and time management. IEEE Transactions on nuclear science, 50(5):1516–1521, 2003.
- [25] Matthew R Palmer and Gordon L Brownell. Annihilation density distribution calculations for medically important positron emitters. IEEE transactions on medical imaging, 11(3):373–378, 1992.
- [26] Craig S Levin and Edward J Hoffman. Calculation of positron range and its effect on the fundamental limit of positron emission tomography system spatial resolution. Physics in Medicine & Biology, 44(3):781, 1999.
- [27] Bing Bai, Ananya Ruangma, Richard Laforest, Y-C Tai, and Richard M Leahy. Positron range modeling for statistical pet image reconstruction. In 2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No. 03CH37515), volume 4, pages 2501–2505. IEEE, 2003.
- [28] Kengo Shibuya, Eiji Yoshida, Fumihiko Nishikido, Toshikazu Suzuki, Tomoaki Tsuda, Naoko Inadama, Taiga Yamaya, and Hideo Murayama. Annihilation photon acollinearity in pet: volunteer and phantom fdg studies. Physics in Medicine & Biology, 52(17):5249, 2007.
- [29] Stefan Seifert, Herman T van Dam, Jan Huizenga, Ruud Vinke, Peter Dendooven, Herbert Löhner, and Dennis R Schaart. Monolithic labr3: Ce crystals on silicon photomultiplier arrays for time-of-flight positron emission tomography. Physics in Medicine & Biology, 57(8):2219, 2012.
- [30] DJ Parker, MR Hawkesworth, CJ Broadbent, P Fowles, TD Fryer, and PA McNeil. Industrial positron-based imaging: principles and applications. Nuclear Instruments and Methods in

CONTENTS

63

- Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 348(2-3):583–592, 1994.
- [31] Kai T Lee, Andy Ingram, and Neil A Rowson. Twin screw wet granulation: the study of a continuous twin screw granulator using positron emission particle tracking (pept) technique. European journal of pharmaceutics and biopharmaceutics, 81(3):666–673, 2012.
- [32] CRK Windows-Yule and DJ Parker. Self-diffusion, local clustering and global segregation in binary granular systems: The role of system geometry. Powder technology, 261:133–142, 2014.
- [33] CRK Windows-Yule and DJ Parker. Boltzmann statistics in a three-dimensional vibrofluidized granular bed: Idealizing the experimental system. Physical Review E, 87(2):022211, 2013.
- [34] Timo Hensler, Martin Tupy, Timo Strer, Thorsten Pöschel, and Karl-Ernst Wirth. Positron emission particle tracking in fluidized beds with secondary gas injection. Powder Technology, 279:113–122, 2015.
- [35] M Stein, TW Martin, JPK Seville, PA McNeil, and DJ Parker. Positron emission particle tracking: particle velocities in gas fluidised beds, mixers and other applications. In Non-invasive monitoring of multiphase flows, pages 309–333. Elsevier, 1997.
- [36] M Stein, YL Ding, JPK Seville, and DJ Parker. Solids motion in bubbling gas fluidised beds. Chemical Engineering Science, 55(22):5291–5300, 2000.
- [37] JA Laverman, X Fan, A Ingram, M van Sint Annaland, DJ Parker, JPK Seville, and JAM Kuipers. Experimental study on the influence of bed material on the scaling of solids circulation patterns in 3d bubbling gas–solid fluidized beds of glass and polyethylene using positron emission particle tracking. Powder Technology, 224:297–305, 2012.
- [38] PK Marsden, RJ Ott, JE Bateman, Simon R Cherry, MA Flower, and S Webb. The performance of a multiwire proportional chamber positron camera for clinical use. Physics in Medicine & Biology, 34(8):1043, 1989.
- [39] MR Hawkesworth, MA O’Dwyer, J Walker, P Fowles, J Heritage, PAE Stewart, RC Witcomb, JE Bateman, JF Connolly, and R Stephenson. A positron camera for industrial application. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 253(1):145–157, 1986.
- [40] CR Bemrose, P Fowles, MR Hawkesworth, and MA O’Dwyer. Application of positron emission tomography to particulate flow measurement in chemical engineering processes. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 273(2-3):874–880, 1988.
- [41] HalO Anger. Gamma-ray and positron scintillation camera. Nucleonics (US) Ceased publication, 21(UCRL-10933), 1963.
- [42] Alireza Sadrumontaz, DJ Parker, and LG Byars. Modification of a medical pet scanner for pept studies. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 573(1-2):91–94, 2007.
- [43] DJ Parker, TW Leadbeater, X Fan, MN Hausard, A Ingram, and Z Yang. Positron emission particle tracking using a modular positron camera. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 604(1-2):339–342, 2009.
- [44] Jonathan Peter Kyle Seville, Yimin Deng, Sarah Dawn Bell, Raf Dewil, Lise Appels, Renaud Ansart, Tom Leadbeater, David Parker, Huili Zhang, Andrew Ingram, et al. 11co2 positron emission imaging reveals the in-situ gas concentration profile as function of time and position in opaque gas-solid contacting systems. Chemical Engineering Journal, 404:126507, 2021.
- [45] TW Leadbeater and DJ Parker. A high speed pc-based data acquisition and control system for positron imaging. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 604(1-2):355–358, 2009.
- [46] Cody Wiggins, Nitant Patel, Zachary Bingham, and Arthur Ruggles. Qualification of multiple-particle positron emission particle tracking (m-pept) technique for measurements in turbulent wall-bounded flow. Chemical Engineering Science, 204:246–256, 2019.

CONTENTS

64

- [47] Nitant Patel, Cody Wiggins, and Arthur Ruggles. Positron emission particle tracking in pulsatile flow. Experiments in Fluids, 58(5):42, 2017.
- [48] YC Tai, A Chatziioannou, S Siegel, J Young, D Newport, RN Goble, RE Nutt, and SR Cherr. Performance evaluation of the micropet p4: a pet system dedicated to animal imaging. Physics in Medicine and Biology, 46:1845–1862, 2001.
- [49] Q Bao, , D Newport, M Chen, DB Stout, and AF Chatziioannou. Performance evaluation of the inveon dedicated pet preclinical tomograph based on the nema nu-4 standards. Journal of Nuclear Medicine, 50:401–408, 2009.
- [50] Z Sarnyai, K Nagy, G Patay, M Molnár, G Rosenqvist, M Tóth, A Takano, B Gulyás, P Major, C Hallden, and A Varrone. Performance evaluation of a high-resolution nonhuman primate pet/ct system. Journal of Nuclear Medicine, 60(12):1818, 2019.
- [51] ME Casey and R Nutt. A multicrystal two dimensional bgo detector system for positron emission tomography. Ieee transactions on nuclear science, 33(1):460–463, 1986.
- [52] TW Leadbeater, DJ Parker, and J Gargiuli. Characterization of the latest birmingham modular positron camera. Measurement Science and Technology, 22(10):104017, 2011.
- [53] TW Leadbeater and DJ Parker. A modular positron camera for the study of industrial processes. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 652(1):646–649, 2011.
- [54] JM Sovechles, D Boucher, R Pax, T Leadbeater, AP Sasmito, and KE Waters. Performance analysis of a new positron camera geometry for high speed, fine particle tracking. Measurement Science and Technology, 28(9):095402, 2017.
- [55] DJ Burnard, AJ Caden, J Gargiuli, T Leadbeater, DJ Parker, and William D Griffiths. A positron emission particle tracking (pept) study of inclusions in liquid aluminium alloy. In Advanced Materials Research, volume 922, pages 43–48. Trans Tech Publ, 2014.
- [56] X Fan, DJ Parker, and MD Smith. Labelling a single particle for positron emission particle tracking using direct activation and ion-exchange techniques. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 562(1):345–350, 2006.
- [57] X. Fan, D.J. Parker, and M.D. Smith. Enhancing ^{18}F uptake in a single particle for positron emission particle tracking through modification of solid surface chemistry. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 558(2):542–546, mar 2006.
- [58] Anthony D Rosato and Christopher Windows-Yule. Segregation in Vibrated Granular Systems. Academic Press, 2020.
- [59] Juan Pellico, Peter J Gawne, and Rafael TM de Rosales. Radiolabelling of nanomaterials for medical imaging and therapy. Chemical Society Reviews, 2021.
- [60] David J Parker and Xianfeng Fan. Positron emission particle tracking—application and labelling techniques. Particuology, 6(1):16–23, 2008.
- [61] Xianfeng Fan, David J Parker, Mike D Smith, Andy Ingram, Zhufang Yang, and Jonathan PK Seville. A simple and selective method for the separation of cu radioisotopes from nickel. Nuclear medicine and biology, 33(7):939–944, 2006.
- [62] Philip W Miller, Nicholas J Long, Ramon Vilar, and Antony D Gee. Synthesis of ^{11}c , ^{18}f , ^{15}o , and ^{13}n radiolabels for positron emission tomography. Angewandte Chemie International Edition, 47(47):8998–9033, 2008.
- [63] François de Dardel and Thomas V. Arden. Ion Exchangers. American Cancer Society, 2008.
- [64] Leila Farzin, Shahab Sheibani, Mohammad Esmaeil Moassesi, and Mojtaba Shamsipur. An overview of nanoscale radionuclides and radiolabeled nanomaterials commonly used for nuclear molecular imaging and therapeutic functions. Journal of Biomedical Materials Research Part A, 107(1):251–285, 2019.
- [65] Maria Grazia Pia, Tullio Basaglia, Zane W Bell, and Paul V Dressendorfer. Geant4 in scientific literature. In 2009 IEEE nuclear science symposium conference record (NSS/MIC), pages

CONTENTS

65

- 189–194. IEEE, 2009.
- [66] J Dudouet, D Cussol, D Durand, and M Labalme. Benchmarking geant4 nuclear models for hadron therapy with 95 mev/nucleon carbon ions. *Physical Review C*, 89(5):054616, 2014.
- [67] Sebastien Incerti, Michael Douglass, Scott Penfold, Susanna Guatelli, and Eva Bezak. Review of geant4-dna applications for micro and nanoscale simulations. *Physica Medica*, 32(10):1187–1200, 2016.
- [68] Sebastien Incerti, Jeremy MC Brown, and Susanna Guatelli. Advances in geant4 applications in medicine. *Physica Medica: European Journal of Medical Physics*, 70:224–227, 2020.
- [69] I Govender, PW Cleary, and AN Mainza. Comparisons of pept derived charge features in wet milling environments with a friction-adjusted dem model. *Chemical Engineering Science*, 97:162–175, 2013.
- [70] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P. M. Bloomfield, D. Brasse, V. Breton, P. Bruyndonckx, I. Buvat, A. F. Chatziioannou, Y. Choi, Y. H. Chung, C. Comtat, D. Donnarieix, L. Ferrer, S. J. Glick, C. J. Groiselle, D. Guez, P.-F. Honore, S. Kerhoas-Cavata, A. S. Kirov, V. Kohli, M. Koole, M. Krieguer, D. J. van der Laan, F. Lamare, G. LARGERON, C. Lartizien, D. Lazaro, M. C. Maas, L. Maigne, F. Mayet, F. Melot, C. Merheb, E. Pennacchio, J. Perez, U. Pietrzyk, F. R. Rannou, M. Rey, D. R. Schaart, C. R. Schmidlein, L. Simon, T. Y. Song, J.-M. Vieira, D. Visvikis, R. Van de Walle, E. Wieërs, and C. Morel. GATE - Geant4 Application for Tomographic Emission: a simulation toolkit for PET and SPECT. *Physics in Medicine and Biology*, 49(19):4543–4561, October 2004.
- [71] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytráček, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell’Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J.J. Gómez Cadenas, I. González, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F.W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampén, V. Lara, V. Lefebvre, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatsu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O’Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M.G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. Di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J.P. Wellisch, T. Wenaus, D.C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschiesche. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.
- [72] Sophie Kerhoas-Cavata and David Guez. Modeling electronic processing in GATE. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 569(2):330–334, December 2006.
- [73] Shoaib Usman and Amol Patil. Radiation detector deadtime and pile up: A review of the status of science. *Nuclear Engineering and Technology*, 50(7):1006–1016, October 2018.
- [74] CRK Windows-Yule, Deepak Raju Tunuguntla, and DJ Parker. Numerical modelling of granular flows: a reality check. *Computational particle mechanics*, 3(3):311–332, 2016.
- [75] Performance Measurements of Positron Emission Tomographs. Technical Report NU-2-2007, National Electronics Manufacturer Association, 2007.
- [76] S. Jan, C. Comtat, D. Strul, G. Santin, and R. Trebossen. Monte carlo simulation for the ecatt

CONTENTS

66

- exact hr+ system using gate. IEEE Transactions on Nuclear Science, 52(3):627–633, Jun 2005.
- [77] John David Anderson and J Wendt. Computational fluid dynamics, volume 206. Springer, 1995.
- [78] Conall J Garvey and Rebecca Hanlon. Computed tomography in clinical practice. BMJ: British Medical Journal, 324(7345):1077–1080, May 2002.
- [79] Hsin-Hon Lin, Keh-Shih Chuang, Yi-Hsing Lin, Yu-Ching Ni, Jay Wu, and Meei-Ling Jan. Efficient simulation of voxelized phantom in gate with embedded simset multiple photon history generator. Physics in Medicine and Biology, 59(20):6231–6250, Oct 2014.
- [80] DJ Parker, CJ Broadbent, P Fowles, MR Hawkesworth, and P McNeil. Positron emission particle tracking—a technique for studying flow within engineering equipment. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 326(3):592–607, 1993.
- [81] DJ Parker. Positron emission particle tracking and its application to granular media. Review of Scientific Instruments, 88(5):051803, 2017.
- [82] Z Yang, DJ Parker, PJ Fryer, S Bakalis, and X Fan. Multiple-particle tracking—an improvement for positron particle tracking. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 564(1):332–338, 2006.
- [83] Z Yang, PJ Fryer, S Bakalis, X Fan, DJ Parker, and JPK Seville. An improved algorithm for tracking multiple, freely moving particles in a positron emission particle tracking system. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 577(3):585–594, 2007.
- [84] M Bickell, A Buffer, I Govender, and DJ Parker. A new line density tracking algorithm for pept and its application to multiple tracers. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 682:36–41, 2012.
- [85] Cody Wiggins, Roque Santos, and Arthur Ruggles. A novel clustering approach to positron emission particle tracking. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 811:18–24, 2016.
- [86] Leonard Kaufmann. Clustering by means of medoids. In Proc. Statistical Data Analysis Based on the L1 Norm Conference, Neuchatel, 1987, pages 405–416, 1987.
- [87] G. Hamerly and C. Elkan. Learning the k in k-means. In Proc. of Advances in Neural Information Processing Systems 16, 2003.
- [88] Charles Romesburg. Cluster analysis for researchers. Lulu. com, 2004.
- [89] Khaled Alsabti, Sanjay Ranka, and Vineet Singh. An efficient k-means clustering algorithm. 1997.
- [90] Ozcan Gundogdu. Multiple particle tracking using the Birmingham positron emission camera. PhD thesis, University of Birmingham, 1999.
- [91] O Gundogdu and E Tarcan. Location-allocation algorithm for multiple particle tracking using birmingham mwpc positron camera. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 523(1-2):223–233, 2004.
- [92] O Gundogdu. Positron emission tomography particle tracking using cluster analysis. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 534(3):562–576, 2004.
- [93] T. Anderson and D. Darling. Asymptotic theory of certain 'goodness of fit' criteria based on stochastic processes. Annals of Mathematical Statistics, 23:193, 1952.
- [94] Cody S Wiggins. Multiple-Particle Positron Emission Particle Tracking (M-PEPT) and its Application to Flows in Porous Media. PhD thesis, University of Tennessee, Knoxville, 2019.
- [95] Kyung Oh Jung, Tae Jin Kim, Jung Ho Yu, Siyeon Rhee, Wei Zhao, Byunghang Ha, Kristy Red-Horse, Sanjiv Sam Gambhir, and Guillem Pratx. Whole-body tracking of single cells via positron emission tomography. Nature biomedical engineering, 4(8):835–844, 2020.
- [96] Gary D Knott. Interpolating cubic splines, volume 18. Springer Science & Business Media, 2000.
- [97] Cody Wiggins, Roque Santos, and Arthur Ruggles. A feature point identification method for

CONTENTS

67

- positron emission particle tracking with multiple tracers. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 843:22–28, 2017.
- [98] I.F. Sbalzarini and P. Koumoutsakos. Feature point tracking and trajectory analysis for video imaging in cell biology. Journal of Structural Biology, 151:182, 2005.
- [99] J.C. Crocker and D.G. Grier. Methods of digital video microscopy for col- loidal studies. Journal of Colloid and Interface Science, 179:298, 1996.
- [100] C.E. Willert and M. Gharib. Digital particle image velocimetry. Experiments in Fluids, 10:181–193, 1991.
- [101] N. Mordant, A.M. Crawford, and E. Bodenschatz. Experimental lagrangian acceleration probability density function measurement. Physica D, 193:245–251, 2004.
- [102] DM Blakemore, I Govender, AT McBride, and AN Mainza. Multiple particle tracking in pept using voronoi tessellations. Chemical Engineering Science, 207:780–789, 2019.
- [103] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. SIAM review, 41(4):637–676, 1999.
- [104] Franz Aurenhammer and Rolf Klein. Voronoi diagrams. Handbook of computational geometry, 5(10):201–290, 2000.
- [105] Franz Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. ACM Computing Surveys (CSUR), 23(3):345–405, 1991.
- [106] Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. Lof: identifying density-based local outliers. In Proceedings of the 2000 ACM SIGMOD international conference on Management of data, pages 93–104, 2000.
- [107] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In Kdd, volume 96, pages 226–231, 1996.
- [108] Kamran Khan, Saif Ur Rehman, Kamran Aziz, Simon Fong, and Sababady Sarasvady. Dbscan: Past, present and future. In The fifth international conference on the applications of digital information and web technologies (ICADIWT 2014), pages 232–238. IEEE, 2014.
- [109] Carine Hue, J-P Le Cadre, and Patrick Pérez. Sequential monte carlo methods for multiple target tracking and data fusion. IEEE Transactions on signal processing, 50(2):309–325, 2002.
- [110] Songhwai Oh, Stuart Russell, and Shankar Sastry. Markov chain monte carlo data association for general multiple-target tracking problems. In 2004 43rd IEEE Conference on Decision and Control (CDC)(IEEE Cat. No. 04CH37601), volume 1, pages 735–742. IEEE, 2004.
- [111] Nicholas T Ouellette, Haitao Xu, and Eberhard Bodenschatz. A quantitative study of three-dimensional lagrangian particle tracking algorithms. Experiments in Fluids, 40(2):301–313, 2006.
- [112] Simo Särkkä, Aki Vehtari, and Jouko Lampinen. Rao-blackwellized particle filter for multiple target tracking. Information Fusion, 8(1):2–15, 2007.
- [113] Samuel S Blackman. Multiple hypothesis tracking for multiple target tracking. IEEE Aerospace and Electronic Systems Magazine, 19(1):5–18, 2004.
- [114] Ayodele E Odo, Indresan Govender, Andy Buffler, and Jean-Paul Franzidis. A pept algorithm for predefined positions of radioisotopes relative to the tracer particle. Applied Radiation and Isotopes, 2019.
- [115] Ricardo JGB Campello, Davoud Moulavi, and Jörg Sander. Density-based clustering based on hierarchical density estimates. In Pacific-Asia conference on knowledge discovery and data mining, pages 160–172. Springer, 2013.
- [116] Sam Manger, Antoine Renaud, and Jacques Vanneste. An expectation–maximization algorithm for positron emission particle tracking. Review of Scientific Instruments, 92(8):085102, 2021.
- [117] Lucien Le Cam. Maximum likelihood: an introduction. International Statistical Review/Revue Internationale de Statistique, pages 153–171, 1990.
- [118] Richard J Rossi. Mathematical statistics: an introduction to likelihood based inference. John

CONTENTS

68

- Wiley & Sons, 2018.
- [119] Douglas A Reynolds. Gaussian mixture models. Encyclopedia of biometrics, 741:659–663, 2009.
- [120] AP Dempster, NM Laird, and DB Rubin. Maximum likelihood from incomplete data via the em algorithm. Journal of the Royal Statistical Society: Series B (Methodological), 39:1–38, 1977.
- [121] Klaus Wienhard, Magnus Dahlbom, Lars Eriksson, Christian Michel, Thomas Bruckbauer, Uwe Pietrzyk, and Wolf-Dieter Heiss. The ecat exact hr: performance of a new high resolution positron scanner. Journal of computer assisted tomography, 18(1):110–118, 1994.
- [122] N Karakatsanis, N Sakellios, NX Tsantilas, N Dikaios, C Tsoumpas, D Lazaro, G Loudos, CR Schmidlein, K Louizi, J Valais, et al. Comparative evaluation of two commercial pet scanners, ecat exact hr+ and biograph 2, using gate. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 569(2):368–372, 2006.
- [123] J Westerweel. Fundamentals of digital particle image velocimetry. Measurement Science and Technology, 8:1379, 1997.
- [124] M Barigou. Particle tracking in opaque mixing systems: an overview of the capabilities of pet and pept. Chemical Engineering Research and Design, 82(9):1258–1267, 2004.
- [125] Thomas W Leadbeater, David J Parker, and Joseph Gargiuli. Positron imaging systems for studying particulate, granular and multiphase flows. Particuology, 10(2):146–153, 2012.
- [126] J Bridgwater, CJ Broadbent, and DJ Parker. Study of the influence of blade speed on the performance of a powder mixer using positron emission particle tracking. Chemical engineering research & design, 71(5):675–681, 1993.
- [127] CJ Broadbent, J Bridgwater, DJ Parker, ST Keningley, and P Knight. A phenomenological study of a batch mixer using a positron camera. Powder Technology, 76(3):317–329, 1993.
- [128] CJ Broadbent, J Bridgwater, and DJ Parker. The effect of fill level on powder mixer performance using a positron camera. The Chemical Engineering Journal and the Biochemical Engineering Journal, 56(3):119–125, 1995.
- [129] JR Jones and J Bridgwater. A case study of particle mixing in a ploughshare mixer using positron emission particle tracking. International Journal of Mineral Processing, 53(1-2):29–38, 1998.
- [130] M Marigo, DL Cairns, M Davies, M Cook, A Ingram, and EH Stitt. Developing mechanistic understanding of granular behaviour in complex moving geometry using the discrete element method. part a: measurement and reconstruction of turbula mixer motion using positron emission particle tracking. Computer Modeling in Engineering & Sciences (CMES), 59(3):217–238, 2010.
- [131] Olga Mihailova, Victor Lim, Michael J McCarthy, Kathryn L McCarthy, and Serafim Bakalis. Laminar mixing in a smx static mixer evaluated by positron emission particle tracking (pept) and magnetic resonance imaging (mri). Chemical Engineering Science, 137:1014–1023, 2015.
- [132] CRK Windows-Yule and DJ Parker. Inelasticity-induced segregation: Why it matters, when it matters. EPL (Europhysics Letters), 106(6):64003, 2014.
- [133] CRK Windows-Yule, BJ Scheper, AJ van der Horn, N Hainsworth, J Saunders, DJ Parker, and AR Thornton. Understanding and exploiting competing segregation mechanisms in horizontally rotated granular media. New journal of physics, 18(2):023013, 2016.
- [134] RD Wildman, S Blackburn, DM Benton, PA McNeil, and DJ Parker. Investigation of paste flow using positron emission particle tracking. Powder technology, 103(3):220–229, 1999.
- [135] J Diemer, C Chilles, J Colbert, T Miri, A Ingram, Philippe David, A Sarhangi Fard, and Patrick D Anderson. Flow visualisation in co-rotating twin screw extruders: positron emission particle tracking and numerical particle trajectories. International Polymer Processing, 26(5):540–550, 2011.
- [136] CRK Windows-Yule, S Gibson, DJ Parker, TZ Kokalova, and JPK Seville. Effect of distributor design on particle distribution in a binary fluidised bed. Powder Technology, 2020.
- [137] Chian W Chan, Jonathan PK Seville, Xianfeng Fan, and Jan Baeyens. Particle motion in cfb cyclones as observed by positron emission particle tracking. Industrial & Engineering

CONTENTS

69

- Chemistry Research, 48(1):253–261, 2009.
- [138] Sarah Forrest, John Bridgwater, Paul R Mort, James Litster, and David J Parker. Flow patterns in granulating systems. Powder Technology, 130(1-3):91–96, 2003.
- [139] Yoshitoshi Saito, Xianfeng Fan, Andy Ingram, and Jonathan Peter Kyle Seville. A new approach to high-shear mixer granulation using positron emission particle tracking. Chemical engineering science, 66(4):563–569, 2011.
- [140] CY Wu, XF Fan, F Motazedian, JPK Seville, DJ Parker, and ACF Cocks. Quantitative investigation of powder flow during die filling using positron emission particle tracking. Proceedings of the Institution of Mechanical Engineers, Part E: Journal of Process Mechanical Engineering, 224(3):169–175, 2010.
- [141] CY Wu, XF Fan, F Motazedian, JPK Seville, DJ Parker, and ACF Cocks. Modelling 1: An experimental study of die filling using positron emission particle tracking. In European Congress and Exhibition on Powder Metallurgy. European PM Conference Proceedings, volume 3, page 335. The European Powder Metallurgy Association, 2007.
- [142] Y-F Chang, CG Ilea, ØL Aasen, and AC Hoffmann. Particle flow in a hydrocyclone investigated by positron emission particle tracking. Chemical Engineering Science, 66(18):4203–4211, 2011.
- [143] Yu-Fen Chang and Alex C Hoffmann. A lagrangian study of liquid flow in a reverse-flow hydrocyclone using positron emission particle tracking. Experiments in Fluids, 56(1):4, 2015.
- [144] Jennifer Rachel Radman, Raymond Langlois, Thomas Leadbeater, James Finch, Neil Rowson, and Kristian Waters. Particle flow visualization in quartz slurry inside a hydrocyclone using the positron emission particle tracking technique. Minerals Engineering, 62:142–145, 2014.
- [145] Rebecca Clare Sindall. Increasing the efficiency of anaerobic waste digesters by optimising flow patterns to enhance biogas production. PhD thesis, University of Birmingham, 2015.
- [146] TS Volkwyn, A Buffler, I Govender, J-P Franzidis, AJ Morrison, A Odo, NP Van Der Meulen, and C Vermeulen. Studies of the effect of tracer activity on time-averaged positron emission particle tracking measurements on tumbling mills at pept cape town. Minerals Engineering, 24(3-4):261–266, 2011.
- [147] Eric N Landis and Denis T Keane. X-ray microtomography. Materials characterization, 61(12):1305–1316, 2010.
- [148] Joshua A Dijkstra, Frank Rietz, Kinga A Lőrincz, Martin van Hecke, and Wolfgang Losert. Invited article: Refractive index matched scanning of dense granular materials. Review of Scientific Instruments, 83(1):011301, 2012.
- [149] Dale L Bailey, Michael N Maisey, David W Townsend, and Peter E Valk. Positron emission tomography, volume 2. Springer, 2005.
- [150] Shantanu Roy, Faical Larachi, MH Al-Dahhan, and MP Duduković. Optimal design of radioactive particle tracking experiments for flow mapping in opaque multiphase reactors. Applied Radiation and Isotopes, 56(3):485–503, 2002.
- [151] Kay A Buist, Alex C van der Gaag, Niels G Deen, and Johannes AM Kuipers. Improved magnetic particle tracking technique in dense gas fluidized beds. AIChE Journal, 60(9):3133–3142, 2014.
- [152] Thorsten Pöschel and Nikolai V Brilliantov. Granular gas dynamics, volume 624. Springer Science & Business Media, 2003.
- [153] JS Olafsen and Jeffery S Urbach. Velocity distributions and density fluctuations in a granular gas. Physical Review E, 60(3):R2468, 1999.
- [154] RD Wildman, JM Huntley, and DJ Parker. Granular temperature profiles in three-dimensional vibrofluidized granular beds. Physical Review E, 63(6):061311, 2001.
- [155] CRK Windows-Yule and DJ Parker. Energy non-equipartition in strongly convective granular systems. The European Physical Journal E, 37(3):1–6, 2014.
- [156] RD Wildman, JM Huntley, J-P Hansen, DJ Parker, and DA Allen. Single-particle motion in three-dimensional vibrofluidized granular beds. Physical Review E, 62(3):3826, 2000.
- [157] RD Wildman, J-P Hansen, and DJ Parker. Velocity auto-correlation functions in three-dimensional vibro-fluidized granular beds. Physics of Fluids, 14(1):232–239, 2002.

CONTENTS

70

- [158] CRK Windows-Yule, AD Rosato, N Rivas, and DJ Parker. Influence of initial conditions on granular dynamics near the jamming transition. *New journal of physics*, 16(6):063016, 2014.
- [159] CRK Windows-Yule, A Moore, C Wellard, D Werner, DJ Parker, and JPK Seville. Particle distributions in binary gas-fluidised beds: Shape matters—but not much. *Chemical Engineering Science*, 216:115440, 2020.
- [160] DJ Parker, AE Dijkstra, TW Martin, and JPK Seville. Positron emission particle tracking studies of spherical particle motion in rotating drums. *Chemical Engineering Science*, 52(13):2011–2022, 1997.
- [161] Gary B Tupper, Indresan Govender, David N De Klerk, Max C Richter, and Aubrey N Mainza. Testing of a new dynamic ergun equation for transport with positron emission particle tracking. *AIChE Journal*, 62(3):939–946, 2016.
- [162] CRK Windows-Yule, AJ Van Der Horn, DR Tunuguntla, DJ Parker, and AR Thornton. Inducing axial banding in bidisperse-by-density granular systems using noncylindrical tumbler geometries. *Physical review applied*, 8(2):024010, 2017.
- [163] P Denissenko, E Guyez, PJ Thomas, DJ Parker, and JPK Seville. Positron emission tracking of individual particles in particle-laden rimming flow. *Physics of Fluids*, 26(5):053304, 2014.
- [164] J Conway-Baker, RW Barley, Richard A Williams, Xiaodong Jia, J Kostuch, B McLoughlin, and DJ Parker. Measurement of the motion of grinding media in a vertically stirred mill using positron emission particle tracking (pept). *Minerals engineering*, 15(1-2):53–59, 2002.
- [165] Rob Fishwick, Mike Winterbottom, David Parker, Xianfeng Fan, and Hugh Stitt. The use of positron emission particle tracking in the study of multiphase stirred tank reactor hydrodynamics. *The Canadian Journal of Chemical Engineering*, 83(1):97–103, 2005.
- [166] Christopher RK Windows-Yule, Roberto Hart-Villamil, Thomas Ridout, Tzany Kokalova, and Jose C Nogueira-Filho. Positron emission particle tracking for liquid-solid mixing in stirred tanks. *Chemical Engineering & Technology*, 43(10):1939–1950, 2020.
- [167] Antonio Guida, Alvin W Nienow, and Mostafa Barigou. Pept measurements of solid–liquid flow field and spatial phase distribution in concentrated monodisperse stirred suspensions. *Chemical Engineering Science*, 65(6):1905–1914, 2010.
- [168] HP Kuo, PC Knight, DJ Parker, Y Tsuji, MJ Adams, and JPK Seville. The influence of dem simulation parameters on the particle behaviour in a v-mixer. *Chemical engineering science*, 57(17):3621–3638, 2002.
- [169] S González, CRK Windows-Yule, Stefan Luding, DJ Parker, and Anthony Richard Thornton. Forced axial segregation in axially inhomogeneous rotating systems. *Physical Review E*, 92(2):022202, 2015.
- [170] Ali Hassanpour, Hongsing Tan, Andrew Bayly, Prasad Gopalkrishnan, Boonho Ng, and Mojtaba Ghadiri. Analysis of particle motion in a paddle mixer using discrete element method (dem). *Powder Technology*, 206(1-2):189–194, 2011.
- [171] HP Kuo, PC Knight, DJ Parker, MJ Adams, and JPK Seville. Discrete element simulations of a high-shear mixer. *Advanced Powder Technology*, 15(3):297–309, 2004.
- [172] M Marigo, M Davies, T Leadbeater, DL Cairns, A Ingram, and EH Stitt. Application of positron emission particle tracking (pept) to validate a discrete element method (dem) model of granular flow and mixing in the turbula mixer. *International journal of pharmaceuticals*, 446(1-2):46–58, 2013.
- [173] TW Martin, JPK Seville, and DJ Parker. A general method for quantifying dispersion in multiscale systems using trajectory analysis. *Chemical engineering science*, 62(13):3419–3428, 2007.
- [174] Yang Yang. *A study of fine particle grinding in vertically stirred media mills via positron emission particle tracking technology and the discrete element method*. PhD thesis, University of Birmingham, 2018.
- [175] Domenico Daraio, Jose Villoria, Andrew Ingram, Alessio Alexiadis, E Hugh Stitt, and Michele Marigo. Validation of a discrete element method (dem) model of the grinding media dynamics

CONTENTS

71

- within an attritor mill using positron emission particle tracking (pept) measurements. Applied Sciences, 9(22):4816, 2019.
- [176] Ali Hassanpour, S Joseph Antony, and Mojtaba Ghadiri. Modeling of agglomerate behavior under shear deformation: effect of velocity field of a high shear mixer granulator on the structure of agglomerates. Advanced Powder Technology, 18(6):803–811, 2007.
- [177] William R Ketterhagen, Mary T am Ende, and Bruno C Hancock. Process modeling in the pharmaceutical industry using the discrete element method. Journal of pharmaceutical sciences, 98(2):442–470, 2009.
- [178] Joel H Ferziger, Milovan Perić, and Robert L Street. Computational methods for fluid dynamics, volume 3. Springer, 2002.
- [179] Christoph Kloss, Christoph Goniva, Alice Hager, Stefan Amberger, and Stefan Pirker. Models, algorithms and validation for opensource dem and cfd-dem. Progress in Computational Fluid Dynamics, an International Journal, 12(2-3):140–152, 2012.
- [180] BGM Van Wachem, JC Schouten, R Krishna, and CM Van den Bleek. Eulerian simulations of bubbling behaviour in gas-solid fluidised beds. Computers & chemical engineering, 22:S299–S306, 1998.
- [181] BGM Van Wachem, JC Schouten, CM Van den Bleek, R Krishna, and JL Sinclair. Cfd modeling of gas-fluidized beds with a bimodal particle mixture. AIChE Journal, 47(6):1292–1302, 2001.
- [182] Jonathan PK Seville and Chuan-Yu Wu. Particle Technology and Engineering: An Engineer's Guide to Particles and Powders: Fundamentals and Computational Approaches. Butterworth-Heinemann, 2016.
- [183] BPB Hoomans, JAM Kuipers, MA Mohd Salleh, M Stein, and JPK Seville. Experimental validation of granular dynamics simulations of gas-fluidised beds with homogenous in-flow conditions using positron emission particle tracking. Powder Technology, 116(2-3):166–177, 2001.
- [184] Yu-Fen Chang, Tom CH Adamsen, Gleb I Pisarev, and Alex C Hoffmann. Pept: An invaluable tool for 3-d particle tracking and cfd simulation verification in hydrocyclone studies. In EPJ Web of Conferences, volume 50, page 05001. EDP Sciences, 2013.
- [185] R Pérez-Mohedano, N Letzelter, C Amador, CT VanderRoest, and S Bakalis. Positron emission particle tracking (pept) for the analysis of water motion in a domestic dishwasher. Chemical Engineering Journal, 259:724–736, 2015.
- [186] José Carlos Nogueira Filho. Effects of mixing regimes and inert solid suspension on sewage sludge anaerobic digestion performance through experimental study and CFD modelling. PhD thesis, University of Birmingham, 2019.
- [187] Fabio Chiti. Lagrangian studies of turbulent mixing in a vessel agitated by a Rushton turbine: positron emission particle tracking (PEPT) and computational fluid dynamics (CFD). PhD thesis, University of Birmingham, 2008.
- [188] Li Liu. Computational fluid dynamics modelling of complex fluid flow in stirred vessels. PhD thesis, University of Birmingham, 2014.
- [189] Paul W Cleary. Dem prediction of industrial and geophysical particle flows. Particuology, 8(2):106–118, 2010.
- [190] YT Feng, Paul W Cleary, Raymond CZ Cohen, Simon M Harrison, Matthew D Sinnott, Mahesh Prakash, and Stuart Mead. Prediction of industrial, biophysical and extreme geophysical flows using particle methods. Engineering Computations, 2013.
- [191] Majid Rasouli. Dynamics of Cylindrical Particles in a Rotating Drum Using Multiple Radioactive Particle Tracking. PhD thesis, École Polytechnique de Montréal, 2015.
- [192] Jose Manuel Valverde and Carlos Soria-Hoyo. Vibration-induced dynamical weakening of pyroclastic flows: Insights from rotating drum experiments. Journal of Geophysical Research: Solid Earth, 120(9):6182–6190, 2015.
- [193] John Mark Nicholas Timm Gray. Granular flow in partially filled slowly rotating drums. Journal of Fluid Mechanics, 441:1–29, 2001.

CONTENTS

72

- [194] D Schneider, R Kaitna, WE Dietrich, L Hsu, C Huggel, and BW McArdell. Frictional behavior of granular gravel–ice mixtures in vertically rotating drum experiments and implications for rock–ice avalanches. *Cold Regions Science and Technology*, 69(1):70–90, 2011.
- [195] Johannes Kulenkampff, Marion Gründig, Abdelhamid Zakhmini, and Johanna Lippmann-Pipke. Geoscientific process monitoring with positron emission tomography (geopet). *Solid Earth*, 7(4):1217–1231, 2016.
- [196] Marion Gründig, Michael Richter, Anita Seese, and Osama Sabri. Tomographic radiotracer studies of the spatial distribution of heterogeneous geochemical transport processes. *Applied Geochemistry*, 22(11):2334–2343, 2007.
- [197] Michael Richter, Marion Gründig, Klaus Zieger, Anita Seese, and Osama Sabri. Positron emission tomography for modelling of geochemical transport processes in clay. *Radiochimica Acta*, 93(9–10):643–651, 2005.
- [198] Qiyong Fan, Akshay Nanduri, Jaewon Yang, Tokihiro Yamamoto, Billy Loo, Edward Graves, Lei Zhu, and Samuel Mazin. Toward a planning scheme for emission guided radiation therapy (egrt): Fdg based tumor tracking in a metastatic breast cancer patient. *Medical physics*, 40(8):081708, 2013.
- [199] Hideaki Tashima, Taiga Yamaya, Eiji Yoshida, Shoko Kinouchi, Mitsuo Watanabe, and Eiichi Tanaka. A single-ring openpet enabling pet imaging during radiotherapy. *Physics in Medicine & Biology*, 57(14):4705, 2012.
- [200] Won Jun Kang, Hyun-Jae Kang, Hyo-Soo Kim, June-Key Chung, Myung Chul Lee, and Dong Soo Lee. Tissue distribution of 18f-fdg-labeled peripheral hematopoietic stem cells after intracoronary administration in patients with myocardial infarction. *Journal of Nuclear Medicine*, 47(8):1295–1301, 2006.
- [201] D Ritchie, L Mileskin, D Wall, J Bartholeyns, M Thompson, J Coverdale, E Lau, J Wong, P Eu, RJ Hicks, et al. In vivo tracking of macrophage activated killer cells to sites of metastatic ovarian carcinoma. *Cancer Immunology, Immunotherapy*, 56(2):155, 2007.
- [202] Xinwei Song, Houston G Wood, and Don Olsen. Computational fluid dynamics (cfd) study of the 4th generation prototype of a continuous flow ventricular assist device (vad). *J. Biomech. Eng.*, 126(2):180–187, 2004.
- [203] D Carswell, D McBride, TN Croft, AK Slone, M Cross, and G Foster. A cfd model for the prediction of haemolysis in micro axial left ventricular assist devices. *Applied Mathematical Modelling*, 37(6):4199–4207, 2013.
- [204] Sadek A Nehmeh and Yusuf E Erdi. Respiratory motion in positron emission tomography/computed tomography: a review. In *Seminars in nuclear medicine*, volume 38, pages 167–176. Elsevier, 2008.
- [205] Keith R Paton, Eswaraiah Varrla, Claudia Backes, Ronan J Smith, Umar Khan, Arlene O’Neill, Conor Boland, Mustafa Lotya, Oana M Istrate, Paul King, et al. Scalable production of large quantities of defect-free few-layer graphene by shear exfoliation in liquids. *Nature materials*, 13(6):624–630, 2014.
- [206] Eswaraiah Varrla, Keith R Paton, Claudia Backes, Andrew Harvey, Ronan J Smith, Joe McCauley, and Jonathan N Coleman. Turbulence-assisted shear exfoliation of graphene using household detergent and a kitchen blender. *Nanoscale*, 6(20):11810–11819, 2014.
- [207] D Loggia, P Gouze, Richard Greswell, and DJ Parker. Investigation of the geometrical dispersion regime in a single fracture using positron emission projection imaging. *Transport in porous media*, 55(1):1–20, 2004.
- [208] Dennis R Schaart. Physics and technology of time-of-flight pet detectors. *Physics in Medicine & Biology*, 2021.
- [209] Sara Reardon. Whole-body pet scanner produces 3d images in seconds. *Nature*, 570(7761):285–287, 2019.