



Identification of the cosmogenic ^{11}C background in large volumes of liquid scintillators with Borexino

M. Agostini^{1,2}, K. Altenmüller², S. Appel², V. Atroshchenko³, Z. Bagdasarian^{4,26}, D. Basilico⁵, G. Bellini⁵, J. Benziger⁶, R. Biondi⁷, D. Bravo^{5,27}, B. Caccianiga⁵, F. Calaprice⁸, A. Caminata⁹, P. Cavalcante^{7,10}, A. Chepurinov¹¹, D. D'Angelo^{5,a}, S. Davini⁹, A. Derbin¹², A. Di Giacinto⁷, V. Di Marcello⁷, X. F. Ding⁸, A. Di Ludovico⁸, L. Di Noto⁹, I. Drachnev¹², A. Formozov^{5,13}, D. Franco¹⁴, C. Galbiati^{8,15}, C. Ghiano⁷, M. Giammarchi⁵, A. Goretti^{8,7}, A. S. Göttel^{4,16}, M. Gromov^{11,13}, D. Guffanti¹⁷, Aldo Ianni⁷, Andrea Ianni⁸, A. Jany¹⁸, D. Jeschke², V. Kobychyev¹⁹, G. Korga^{20,21}, S. Kumaran^{4,16}, M. Laubenstein⁷, E. Litvinovich^{3,22}, P. Lombardi⁵, I. Lomskaya¹², L. Ludhova^{4,16}, G. Lukyanchenko³, L. Lukyanchenko³, I. Machulin^{3,22}, J. Martyn¹⁷, E. Meroni⁵, M. Meyer²³, L. Miramonti⁵, M. Misiaszek¹⁸, V. Muratova¹², B. Neumair², M. Nieslony¹⁷, R. Nugmanov^{3,22}, L. Oberauer², V. Orekhov¹⁷, F. Ortica²⁴, M. Pallavicini⁹, L. Papp², L. Pelicci^{4,16}, Ö. Penek⁴, L. Pietrofaccia⁸, N. Pilipenko¹², A. Pocar²⁵, A. Porcelli^{17,28}, G. Raikov³, M. T. Ranalli⁷, G. Ranucci^{5,b}, A. Razeto⁷, A. Re⁵, M. Redchuk^{4,16,29}, A. Romani²⁴, N. Rossi⁷, S. Schönert², D. Semenov¹², G. Settanta⁴, M. Skorokhvatov^{3,22}, A. Singhal^{4,16}, O. Smirnov¹³, A. Sotnikov¹³, Y. Suvorov^{3,7,30}, R. Tartaglia⁷, G. Testera⁹, J. Thurn²³, E. Unzhakov¹², A. Vishneva¹³, R. B. Vogelaar¹⁰, F. von Feilitzsch², M. Wojcik¹⁸, M. Wurm¹⁷, S. Zavatarelli⁹, K. Zuber²³, G. Zuzel¹⁸

- ¹ Department of Physics and Astronomy, University College London, London, UK
- ² Physik-Department and Excellence Cluster Universe, Technische Universität München, 85748 Garching, Germany
- ³ National Research Centre Kurchatov Institute, 123182 Moscow, Russia
- ⁴ Institute for Nuclear Physics IKP-2, Forschungszentrum Jülich, 52428 Jülich, Germany
- ⁵ Dipartimento di Fisica, Università degli Studi e INFN, 20133 Milan, Italy
- ⁶ Chemical Engineering Department, Princeton University, Princeton, NJ 08544, USA
- ⁷ INFN Laboratori Nazionali del Gran Sasso, 67010 Assergi, AQ, Italy
- ⁸ Physics Department, Princeton University, Princeton, NJ 08544, USA
- ⁹ Dipartimento di Fisica, Università degli Studi e INFN, 16146 Genoa, Italy
- ¹⁰ Physics Department, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061, USA
- ¹¹ Lomonosov Moscow State University Skobel'syn Institute of Nuclear Physics, 119234 Moscow, Russia
- ¹² St. Petersburg Nuclear Physics Institute NRC Kurchatov Institute, 188350 Gatchina, Russia
- ¹³ Joint Institute for Nuclear Research, 141980 Dubna, Russia
- ¹⁴ AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/IRFU, Observatoire de Paris, Sorbonne Paris Cité, 75205 Paris Cedex 13, France
- ¹⁵ Gran Sasso Science Institute (INFN), 67100 L'Aquila, Italy
- ¹⁶ Physics Institute IIIB, RWTH Aachen University, 52062 Aachen, Germany
- ¹⁷ Institute of Physics and Excellence Cluster PRISMA+, Johannes Gutenberg-Universität Mainz, 55099 Mainz, Germany
- ¹⁸ M. Smoluchowski Institute of Physics, Jagiellonian University, 30059 Kraków, Poland
- ¹⁹ Kiev Institute for Nuclear Research, Kyiv 03680, Ukraine
- ²⁰ Department of Physics, Royal Holloway University of London, Egham, Surrey TW20 0EX, UK
- ²¹ Institute of Nuclear Research (Atomki), Debrecen, Hungary
- ²² National Research Nuclear University MEPhI (Moscow Engineering Physics Institute), 115409 Moscow, Russia
- ²³ Department of Physics, Technische Universität Dresden, 01062 Dresden, Germany
- ²⁴ Dipartimento di Chimica, Biologia e Biotechnologie, Università degli Studi e INFN, 06123 Perugia, Italy
- ²⁵ Amherst Center for Fundamental Interactions and Physics Department, University of Massachusetts, Amherst, MA 01003, USA
- ²⁶ Present address: Department of Physics, University of California, Berkeley, Berkeley, CA 94720, USA
- ²⁷ Present address: Universidad Autónoma de Madrid, Ciudad Universitaria de Cantoblanco, 28049 Madrid, Spain
- ²⁸ Present address: Departement Fysica en Sterrenkunde, Universiteit Gent, Gent, Belgium
- ²⁹ Present address: Dipartimento di Fisica e Astronomia dell'Università di Padova and INFN Sezione di Padova, Padua, Italy
- ³⁰ Present address: Dipartimento di Fisica, Università degli Studi Federico II e INFN, 80126 Naples, Italy

Received: 19 June 2021 / Accepted: 2 November 2021
© The Author(s) 2021

Abstract Cosmogenic radio-nuclei are an important source of background for low-energy neutrino experiments. In Borexino, cosmogenic ^{11}C decays outnumber solar *pep* and CNO neutrino events by about ten to one. In order to extract the flux of these two neutrino species, a highly efficient identification of this background is mandatory. We present here the details of the most consolidated strategy, used throughout Borexino solar neutrino measurements. It hinges upon finding the space-time correlations between ^{11}C decays, the preceding parent muons and the accompanying neutrons. This article describes the working principles and evaluates the performance of this Three-Fold Coincidence (TFC) technique in its two current implementations: a hard-cut and a likelihood-based approach. Both show stable performances throughout Borexino Phases II (2012–2016) and III (2016–2020) data sets, with a ^{11}C tagging efficiency of $\sim 90\%$ and $\sim 63\text{--}66\%$ of the exposure surviving the tagging. We present also a novel technique that targets specifically ^{11}C produced in high-multiplicity during major spallation events. Such ^{11}C appear as a *burst* of events, whose space-time correlation can be exploited. Burst identification can be combined with the TFC to obtain about the same tagging efficiency of $\sim 90\%$ but with a higher fraction of the exposure surviving, in the range of $\sim 66\text{--}68\%$.

1 Introduction

Since 2007, the Borexino detector has been observing solar neutrinos from the depths of the Gran Sasso National Laboratory (Italy). Designed for the spectroscopy of low energy neutrinos detected via elastic scattering off electrons, Borexino has accomplished individual measurements of the solar ^7Be [1], ^8B [2], *pep* [3], and *pp* neutrino fluxes [4]. Moreover, Borexino has performed a comprehensive spectroscopic measurement of the entire neutrino spectrum emitted in the *pp*-chain [5–7]. Most recently, Borexino achieved the first direct observation of neutrinos produced in the solar CNO cycle [8,9].

While there are several elements that contributed to the success of Borexino, we focus here on the analysis strategy that has been developed for the identification of cosmogenic background in the 1–2 MeV spectral region (originally proposed in [10]). The observation of *pep* and CNO neutrinos is made possible by the suppression of the β^+ -decays of ^{11}C that are otherwise largely indistinguishable from neutrino-induced electron recoils on an event-by-event basis. In organic scintillators, ^{11}C is being produced by cosmic muons that penetrate the rock coverage of the laboratory and induce spallation processes on carbon of type

$\mu + ^{12}\text{C} \rightarrow \mu + ^{11}\text{C} + n$. A discussion of the processes that contribute to ^{11}C isotope production can be found in [11]. The isotope features a lifetime of 29.4 min and a Q value of 0.96 MeV. Due to positron annihilation, the visible spectrum is shifted to higher energies, covering a range between ~ 0.8 and ~ 2 MeV.

Figure 1 shows the main neutrino signal and radioactive background contributions to the visible energy spectrum of Borexino. At first glance, detection of *pep* and CNO neutrinos seems not very promising since other solar neutrino species cover the lower part of their spectra while towards the upper end ^{11}C overtops their signals by about an order of magnitude. An efficient veto strategy for ^{11}C is a necessary prerequisite for overcoming this disadvantageous situation.

The present article illustrates two analysis techniques for an effective ^{11}C suppression. The first is the Three-Fold Coincidence (TFC) technique. The TFC is formed by the space and time correlation of ^{11}C events with their parent muons and the neutrons that in most cases accompany ^{11}C production. By now, the technique has been extensively used in Borexino analyses. TFC creates two subsets of data: (1) a ^{11}C -depleted set that is including over 60% of the exposure but less than 10% of the ^{11}C events; (2) a ^{11}C -enriched set with the complementary exposure and ^{11}C events. Both datasets are used in the multivariate fit that leads to the extraction of the solar neutrino signals as explained in [6,12].

The second technique is the Burst Identification (BI) technique, a new tagging strategy that targets ^{11}C events originating from a single hadronic shower upon the passage of a muon. These events are visible in *bursts* and are space-time correlated among themselves: due to the negligible convective motion of the scintillator in the timescale of the ^{11}C mean life, they remain essentially aligned along the parent muon track. Since the BI technique on its own does not provide sufficient tagging efficiency for solar neutrino analyses, it is only used in combination with the TFC in Borexino. However, the BI method might be interesting as a stand-alone solution for experimental setups that do not provide sufficient information for a TFC-like approach.

In this article, we first recall the key features of the Borexino detector (Sect. 2); we then review the TFC technique, presenting the two most developed implementations that use either a cut-based or likelihood-based approach and their comparative performances (Sect. 3); we go on to discuss the new BI technique and the gain in ^{11}C -depleted exposure that can result from a combined application of TFC and BI (Sect. 4), before presenting our conclusions in Sect. 5.

It is worth to note that both the TFC and the BI algorithms have been developed for the tagging of ^{11}C but are not limited to this specific cosmogenic isotope. Despite some variability in the underlying production processes [13], the creation of radioactive isotopes in organic scintillator is closely linked to muons inducing hadronic showers, resulting in especially

^a e-mail: davide.dangelo@mi.infn.it (corresponding author)

^b e-mail: spokesperson-borex@lngs.infn.it (corresponding author)

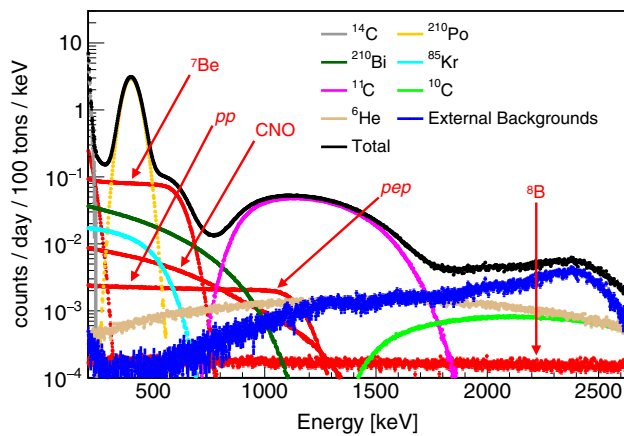


Fig. 1 Expected spectrum in Borexino from Monte Carlo simulations. Electron recoil energy due to neutrino interactions (red lines) and background components (other colors). The Fiducial Volume applied is described in Sect. 2

bright muon events and subsequent *bursts* of delayed neutron captures. Both tagging techniques are of immediate use for upcoming neutrino detectors based on organic scintillators [14, 15] but may as well be of interest for water Cherenkov detectors with neutron-tagging capabilities [16–18] or hybrid detectors [19].

2 The Borexino detector

A schematic drawing of the Borexino detector [20] is shown in Fig. 2. The neutrino target consists of 278 t of organic scintillator composed of the solvent PC (1,2,4-trimethylbenzene) doped with the wavelength shifter PPO (2,5-diphenyloxazole) at a concentration of 1.5 g/l. The scintillator mixture is contained in a spherical and transparent nylon Inner Vessel (IV) with a diameter of 8.5 m and a thickness of 125 μm . To shield this central target from external γ -ray backgrounds and to absorb emanating radon, the IV is surrounded by two layers of buffer liquid in which the light quencher DMP (dimethylphthalate) is added to the scintillator solvent. A Stainless Steel Sphere (SSS) of 13.7 m diameter holding 2212 8" photomultiplier tubes (PMTs) completes the Inner Detector (ID). PMTs are inward-facing and detect the scintillation light caused by particle interactions in the central region. The ID is embedded in a steel dome of 18 m diameter and 16.9 m height that is filled with 2.1 kt of ultra-pure water. The outer surface of the SSS and the floor of the water tank are instrumented with additional 208 PMTs. This Outer Detector (OD) provides efficient detection and tracking of cosmic muons via the Cherenkov light that is emitted during their passage through the water. Neutrons created by the passage of muons are quickly thermalized and captured on hydrogen (carbon in 1% of the cases) in the scintillator after $\tau_n \sim 260 \mu\text{s}$ [21] with the emission of

a 2.2 MeV (4.9 MeV) de-excitation gamma-ray. Muon and neutron detection in Borexino is described in detail in [22]. A thorough analysis of all cosmogenic backgrounds can be found in [21] and the most recent update on the cosmic muons and neutrons in [23].

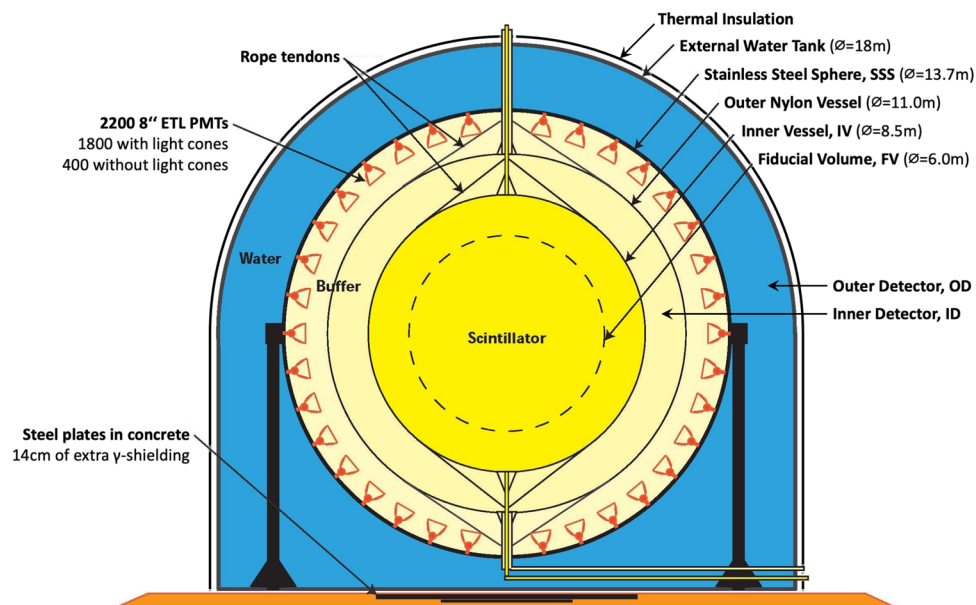
Event reconstruction. In Borexino, a physical event is defined by the firing of several PMTs of the ID within a short period of time (typically 20–25 PMTs in 60–100 ns, changing over time). This occurrence causes a trigger signal and opens an acquisition window of 16 μs during which all PMT pulses are saved. As this window may contain more than one physical event, an offline *clustering* algorithm identifies pulses within a short time window (several 100 ns) as belonging to a specific event. Since the scintillation light output is proportional to the deposited energy, the number of pulses (“hit”) recorded by the PMTs for a given *cluster*, N_{hits} , is used as energy estimator. We record about 500 photoelectrons per MeV of deposited energy. In order to account for the increasing number of malfunctioning PMTs throughout the data set [6], the number of hits is always normalized to 2000 live channels. The position of the scintillation events in the scintillator volume is reconstructed based on the time-of-flight of the scintillation photons from vertex to PMTs. The reconstruction algorithm uses PDFs for the expected arrival time distributions at the PMTs to take into account both propagation effects and the time profile of the scintillation process [12].

Muon and neutron windows. A muon crossing the scintillator target is expected to cause sizable signals in both ID and OD.¹ As for ordinary events, the ID trigger opens a 16 μs -long “Main Window” (MW) for which all hits are saved. The OD issues an additional trigger flag with high efficiency. In case of its presence, the MW is followed by a substantially longer “Neutron Window” (NW): 1.6 ms. This is meant to catch the subsequent signals due to the capture of muon-induced neutrons without requiring explicit triggers. Given the value of τ_n , about 6% of neutrons are collected in the MW and the rest in the NW.

Empty boards condition. The ID is designed for low energy spectroscopy, so scintillator-crossing muons that generate a huge amount of light have a significant chance of saturating the electronics memory. Consequently, the reconstruction of the energy and of the position of the neutron-capture gammas following such event can be troublesome. Effects are more severe in the MW than in the NW. It has been found that this condition can be quantified by the number of electronic boards that show no data for MW and NW in any of their eight channels when memory saturation occurs. Thus, the number of boards devoid of hits (out of a total of 280 boards)

¹ Muons crossing only the OD instead do not play a role in this analysis.

Fig. 2 Sketch of the Borexino detector



is counted and stored for analysis as parameter N_{EB} (# of Empty Boards).

Neutron identification. Luminous muon events produce a long tail of trailing hits that extend into the NW (e.g. by after pulsing of PMTs). A specifically adapted *clustering* algorithm has been developed for both the muon MW and NW to find physical events superimposed to this background. This task is further complicated by the aforementioned presence of saturation effects.

Due to its considerable length the NW contains not only neutrons (and cosmogenic isotopes) but for the most part accidental counts. These are dominated by low energy ^{14}C decays that occur at a rate of about 100 s^{-1} [4]. To avoid mis-counting those and other low-energy events as neutron captures, an energy selection threshold has been introduced at $N_{thr} = 385$ hits (corresponding roughly to 770 keV). To take into account the saturation effects that diminish the number of hits collected from a neutron event, the selection criterium is adjusted to $N_{hits} > N_{thr} - 2 \times N_{EB}$.

Neutron vertex reconstruction. Following bright muons, the accuracy of vertex reconstruction can be severely affected by electronics saturation effects. Therefore, position reconstruction is ineffective for neutrons in the MW and can still be severely affected in the NW. The TFC algorithms described below take this into account by introducing Full Volume Vetoes (see Sect. 3.3) in presence of severe saturation effects. The BI algorithm described in Sect. 4 can be used as an alternative, more effective way to target ^{11}C produced under these circumstances.

Fiducial volume. Unless otherwise noted, in this paper we refer to fiducial volume as to the one used in the *pp*-chain [5] and in the CNO [8] neutrino analyses. This means con-

sidering events with a reconstructed position within a 2.8 m radius from the center of the detector and vertical coordinate in the range $[-1.8, 2.2]$ m.

Data taking campaigns. The Borexino data used throughout this paper refer to Phase-II and Phase-III of the data taking campaign, ranging respectively from 14 December 2011 till 21 May 2016 and from 17 July 2016 till 1 March 2020. The average number of working PMTs during Phase-II and III was about 1500 and about 1200, respectively.

A temperature control strategy was gradually adopted throughout Phase-III to prevent convection in the liquid scintillator [8] in order to stabilize internal radioactive background levels. However, we note that no impact is expected in ^{11}C suppression mechanisms as the observed convective time scale is much longer than ^{11}C mean life.

3 Three-fold coincidence

The TFC identifies ^{11}C candidates by their space and time correlation with the preceding parent muons and muon-induced neutrons and splits the original data sample into two sub-samples: one containing the ^{11}C -tagged events (enriched sample) and another one where most ^{11}C events are removed (depleted sample).

In practice, the procedure is complicated by the relatively long lifetime of the ^{11}C isotope ($\tau_{^{11}\text{C}} = 29.4$ min) compared to the rate of muons crossing the Borexino ID (3 min^{-1}). While the association of muon tracks and subsequent neutron capture vertices ($\tau_n \sim 260\ \mu\text{s}$) is straight-forward, tagging conditions have to be maintained for several $\tau_{^{11}\text{C}}$ and thus select not only ^{11}C candidates but also substantial amounts of non-cosmogenic events. While we refer to the latter as

neutrino events in the following, it is worth to point out that this event category contains as well a large amount of decays caused by internal and external radioactivity that we include under this denomination for the sake of simplicity.

Typically, about 40% of the exposure is contained in the ^{11}C -enriched data set that includes over 90% of ^{11}C decays. The complementary ^{11}C -depleted data set thus mostly consists of *neutrino* events and corresponds to $\sim 60\%$ of the exposure. This latter set is the most relevant for *pep* and CNO neutrino analyses.² Thus, the goal of the TFC algorithms is to minimize the ^{11}C that remains in the ^{11}C -depleted subset and to maximize the fraction of the exposure that contributes to it.

We note that in [11] it is theoretically calculated that about 5% of the ^{11}C production does not have a neutron associated. Although in principle this could imply an inefficiency of the TFC method, several ^{11}C atoms and additional neutrons are generally produced by the same parent muon, effectively limiting the impact of these invisible channels.

This section describes the two TFC approaches used in Borexino analyses: the Hard Cut approach (HC-TFC, Sect. 3.1) and the Likelihood approach (LH-TFC, Sect. 3.2).

HC-TFC generates a list of space-time veto regions within the exposure upon meeting certain conditions (typically a muon followed by one or more neutrons). All regular events are then tested against the regions in the list and if they fall within one (or more) they are labelled as ^{11}C . LH-TFC builds for each event a likelihood to be a ^{11}C decay based of its space-time distance from preceding muon-neutron pairs and applies the ^{11}C tag if the likelihood exceeds a programmable threshold. ^{11}C -tagged events will include a fraction of falsely tagged *neutrino* events that are thus included in the ^{11}C -enriched spectrum. Vice versa, the ^{11}C -depleted data set will contain a residual of ^{11}C decays. Their fraction represents the inefficiency of the TFC method.

Figure 3 shows the ^{11}C -depleted and -enriched visible energy spectra based on Borexino Phase-II data. The effect of the TFC in the ^{11}C energy range (0.75–1.87 MeV) is striking. A difference in the spectral shapes persists at even higher energies. This difference can be attributed to other cosmogenic isotopes (predominantly ^{10}C and ^6He) that feature higher spectral end-points and that are picked up as well by the TFC tagging.

The fraction of exposure that contributes to each spectrum is computed within each TFC algorithm by generating random events with uniform time and position distributions within the fiducial volume, effectively reproducing every real run. These events are then tested against the same list of veto regions or likelihood condition of the real events.

² Nevertheless, Borexino multivariate fit strategy uses both subsets as explained in [6, 12]. This maximises the sensitivity to determine neutrino and background signals outside the ^{11}C energy region.

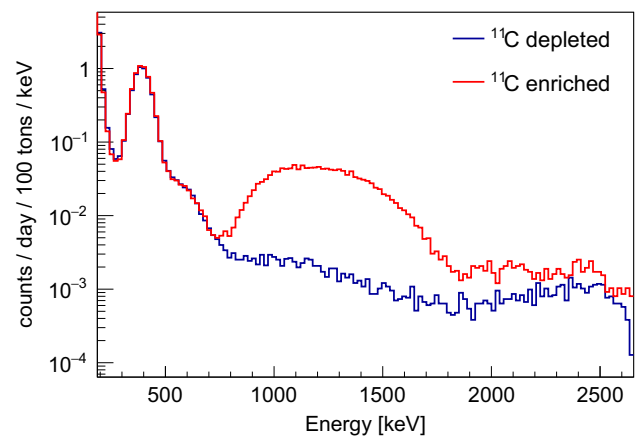


Fig. 3 ^{11}C -depleted (blue line) and enriched (red line) charge spectra created with the TFC method for Phase-II. Spectra are normalized to the same exposure. Only basic cuts that remove muon and noise events are applied. The spectral difference above 1800 keV is caused by the TFC tagging of cosmogenic ^6He and ^{10}C

3.1 Hard cut approach

The first approach for the suppression of the ^{11}C background developed in Borexino [12] is based on defining a list of space-time veto regions in the dataset. A region is added to the list upon the occurrence of a combination of muon and neutron events. Each region holds for a programmable time after the events that generated it. We set this time to 2.5 h, approximately five times the ^{11}C mean time, based on an optimization for tagging efficiency and exposure (see Sect. 3.4). Regular events are then checked against the list of regions and if falling within one or more they are flagged as ^{11}C .

We now describe the different conditions that determine the creation of a veto region. In the definition of the regions we use the concept of neutron multiplicity, indicating how many clusters of PMT hits that follow a muon event could possibly be due to a *n*-capture gamma. In the MW we count $n_{\mu n}$ as the number of clusters found by the ad-hoc cluster-finding algorithm that accounts for the underlying tail of the muon event, excluding the first $3.5\ \mu\text{s}$ of the window. In the NW we count n_n as the number of clusters that meet the requirement expressed in Sect. 2.

Upon finding a muon followed by at least one neutron cluster, we define geometrical veto regions: (1) a cylindrical region around the reconstructed muon track (if available); (2) a spherical region around the *n*-capture reconstructed position in the NW; (3) a spherical region around the projection point of the *n*-capture position in the NW onto the muon track (if available). The working principle is depicted in Fig. 4. After tuning for tagging efficiency and exposure (see Sect. 3.4), the radius of the cylindrical and spherical

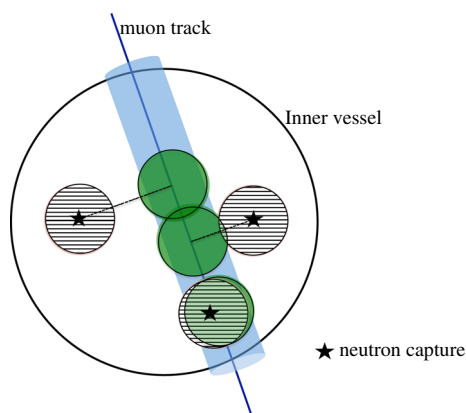


Fig. 4 HC-TFC working principle. Upon the passage of a muon through the IV, geometrical veto regions are defined: a cylinder around the muon track (blue), and spheres around the n-capture points (shaded) and around their projection upon the track (green). The typical radius of the cylindrical (spherical) regions is 0.7 m (1.2 m). The picture is not to scale

regions are chosen to be 0.7 m and 1.2 m, respectively, as summarized in Fig. 4 caption.

Concerning the tracking of muons, as reported in [22], Borexino has developed two algorithms based on the ID signals and one based on the Cherenkov light recorded by the OD. The track resulting from a global fit of four entry/exit-points (two from either of the ID-based algorithms and two from the OD-based one) is chosen for the vast majority of muons. In those cases when the fit result is inadequate, we rely on the available results of the individual algorithms [22]. In the rare cases ($\sim 0.1 \mu/d$) for which none of the tracking algorithms could provide a meaningful track, we apply a Full Volume Veto (introduced in Sect. 3.3).

In the years up to 2012, the CNGS neutrino beam aimed from CERN to Gran Sasso induced ν_μ interactions in both the Borexino detector and the rock upstream, creating a sizable amount of muons crossing the scintillator volume [24]. These muons are relatively low in energy (~ 20 GeV) and, as expected, we observe that they do not significantly contribute to ^{11}C background. Therefore, we exclude them from the potentially veto-triggering muons based on the beam time stamps.

In addition to the veto regions described in this section, HC-TFC also implements the Full Volume Vetoes described in Sect. 3.3.

3.2 Likelihood approach

The implementation of the TFC described above applies hard cuts between muons, neutrons, and β -events to make a binary decision whether a candidate event should be regarded as a *neutrino* or ^{11}C event. In contrast, the approach detailed in

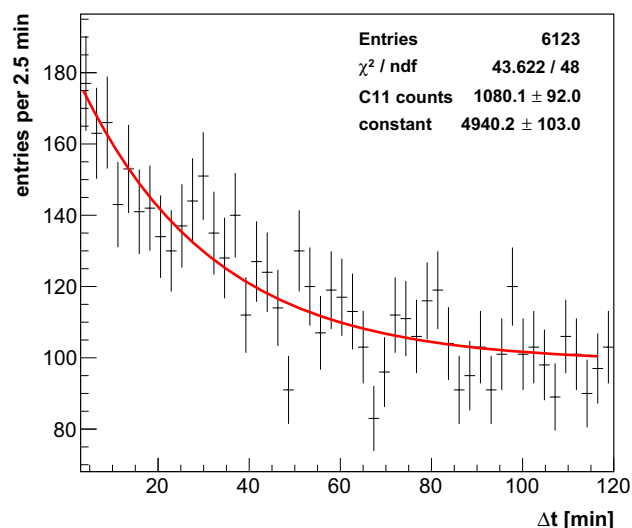


Fig. 5 Example of a Δt profile of ^{11}C -candidates following muons. The distribution contains an exponentially decaying component from correlated muon- ^{11}C pairs as well as the flat component from accidental coincidences. In this special case, a lateral distance cut of $d_\mu = [1.0; 1.2]$ m is applied. The ratio of correlated to uncorrelated pairs is used to produce the probability profiles for the LH-TFC tagging shown in Fig. 6; in this case, the 6th bin of histogram (b)

this section tries a further optimization of tagging efficiency and dead-time reduction by assigning each event a non-binary likelihood to be cosmogenic. A likelihood $\mathcal{L} = \prod_i p_i$ is obtained from the product of several partial probabilities to be ^{11}C , p_i , that are based on a set of TFC discrimination parameters.

The discrimination parameters are largely the same as for the HC-TFC (Sect. 3.1), i.e. the spatial distances and time elapsed with regards to the potential parent muons and neutrons. In addition, it permits to directly include the neutron multiplicity, the presence of neutrons within the MW, and the differential energy loss dE/dx of the parent muon. The corresponding probability profiles p_i can be obtained directly from data (see below).

The rationale behind this approach is that a ^{11}C candidate closely missing all hard-cut criteria to be cosmogenic would still score high in the likelihood approach, while an event fulfilling only narrowly one selection criterium will not be selected. Thus, the likelihood is expected to do particularly well in those cases. It also permits to balance between spatial and time distance between muons (or neutrons) and ^{11}C candidates, meaning effectively that wider spatial cuts are implemented in case of shorter time differences.

A similar tagging scheme for the cosmogenic isotope ^9Li has been successfully applied in the Double-Chooz experiment [25, 26].

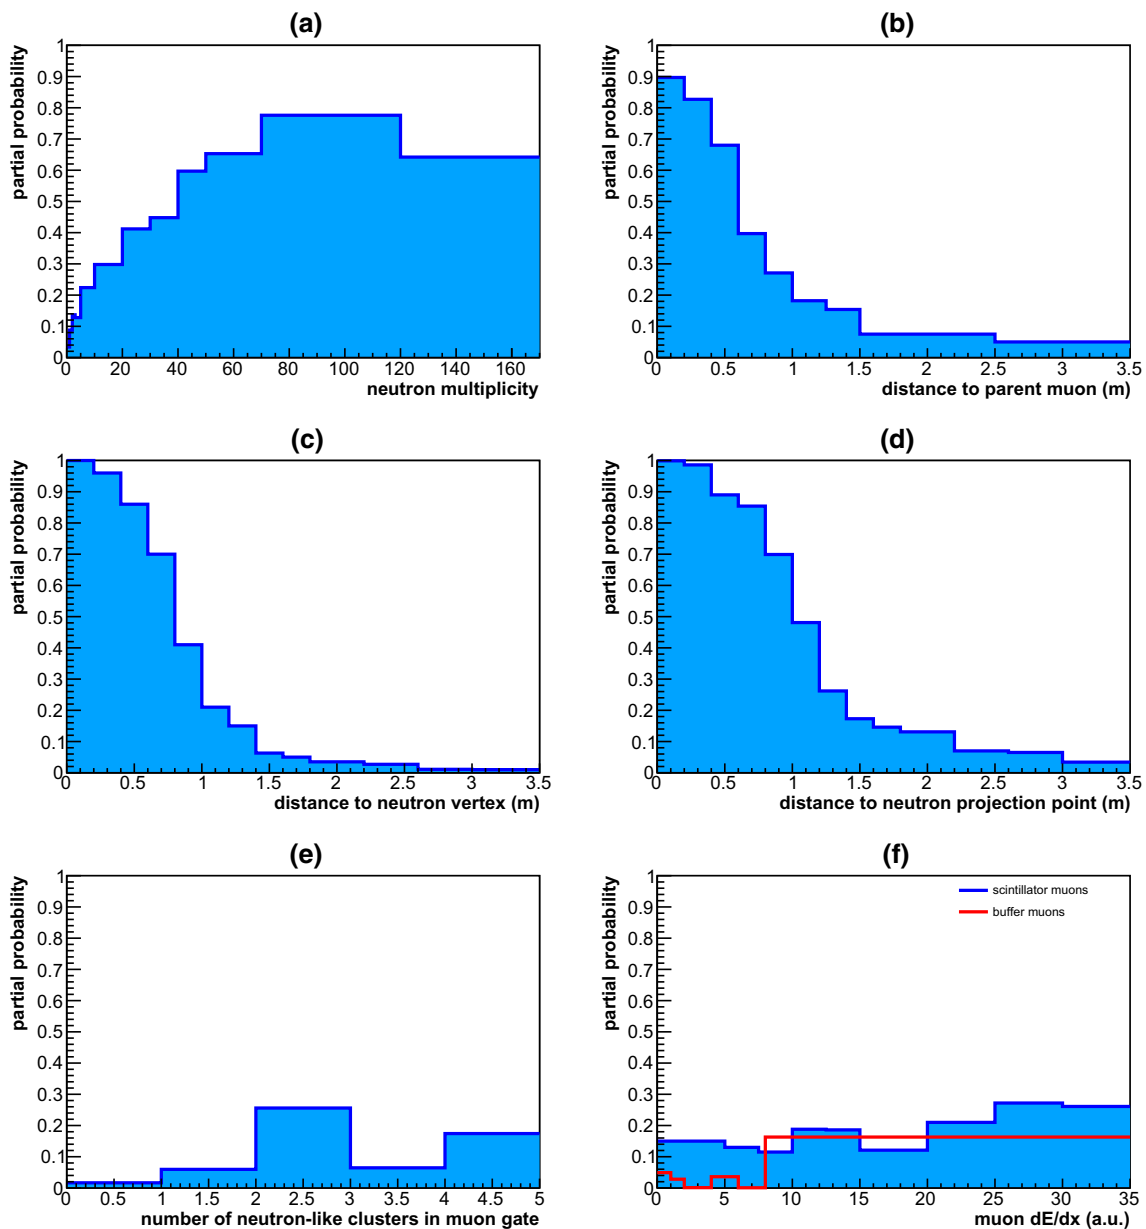


Fig. 6 Overview of partial probability profiles of likelihood TFC discrimination parameters (description in the text)

3.2.1 Data-driven probability profiles

The majority of probability profiles has been derived directly from the data, using the 2012 data set. For a given discrimination parameter i , the partial probability p_i is determined by sorting the data set into several subsamples based on intervals of the parameter value. In each subsample, the relative contributions of ^{11}C and uncorrelated events is evaluated based on the time difference (Δt_k) profile between a ^{11}C -candidate and all preceding muons (k) in the last two hours. As an example, Fig. 5 shows the Δt_k profile of ^{11}C -candidates selected by the distance from the parent muon track d_μ in the interval

[1.0 m; 1.2 m). The Δt_k distribution contains a flat component associated to accidental coincidences of unrelated μ -candidate pairs as well as an exponentially decaying component (following the ^{11}C lifetime) of true μ - ^{11}C coincidences. As shown in Fig. 5, the relative fraction of correlated ^{11}C in the subsample can be extracted by a fit containing those two components. This fraction is used as probability value p_i in the construction of the partial probability profile. For each bin in a given profile, the corresponding data sample is selected, the Δt_k distribution is fitted and the exponential fraction is determined.

We use a selection of seven discrimination parameters. The probability profiles obtained from the data for the first six are displayed in Fig. 6:

(a) **Neutron multiplicity**, $p(n_n)$: A higher number of neutrons, n_n , following a parent muon is a good indicator for hadronic shower formation. In some cases, a very luminous muon will flood the internal memories of the electronics, causing the NW to be entirely devoid of hits; based on a fit of the corresponding Δt -distribution (i.e. selecting only muons followed by no hits in the subsequent NW), we determine for this case a value of $p(n_n) = 0.885$, corresponding to a very high probability to find subsequent ^{11}C decays.

(b) **Distance to the muon track**, $p(d_\mu)$: The probability to be ^{11}C is found to decline as function of d_μ . In cases where muon track reconstruction fails, a probability $p(d_\mu)$ of 0.182 is assigned (based on the corresponding subsample).

(c) **Distance to the closest neutron vertex**, $p(d_n)$: Proximity to the neutron capture vertex was found to be a strong indicator for ^{11}C .

(d) **Distance to the closest neutron projection point**, $p(d_{np})$: Proximity to the projection of the vertex on the muon track was identified as a similarly efficient indicator.

(e) **Neutrons in the MW**, $p(n_{\mu n})$: While the entire MW in the wake of the initial muon signal is usually dominated by afterpulses, a special clustering algorithm has been devised to find superimposed peaks from cosmogenic isotopes and neutrons. While the vertex of these events cannot be reconstructed, the mere presence of these late clusters provides a (weak) indication of shower formation and hence ^{11}C production.

(f) **dE/dx of the parent muon**, $p(dE/dx)$: High energy deposition per unit muon path length is a good indicator for a muon-induced hadronic shower. Here, we consider the visible energy of the muon per total path length as an observational estimator for the physical quantity. Two different distributions are used, one describing muons that have crossed the scintillator volume (scintillator muons), the other for those passing it by (buffer muons).

(g) **Time elapsed since muon**, $p(\Delta t)$: The radioactive decay of ^{11}C means that a large time difference to the parent muon are less likely. The basic probability assigned is an exponential decay with the ^{11}C lifetime.

It is worthwhile to note that the partial profiles obtained do not constitute probability density functions in a mathematical sense. For instance, a good fraction of the accidental coincidences has been removed from the Δt profiles by selecting only events inside the ^{11}C energy range and within a spherical volume of 3 m radius to remove both low-energy radioactive

decays and external backgrounds. For some of the variables that only weakly select for ^{11}C (e.g. the muon dE/dx), it was necessary to apply additional cuts to a second discrimination parameter to further enhance the signal-to-background ratio. Finally, the integral of the distributions shown in Fig. 6 are not normalized. However, this is of no practical consequence for the construction of the likelihood since the contribution of the p_i are a posteriori tuned by weight factors w_i (see below).

3.2.2 Construction of the likelihood

The construction of a single likelihood \mathcal{L} for a given ^{11}C candidate necessitates the expansion of the simple likelihood definition to the expression

$$\mathcal{L}_{jk} = q_j \cdot \prod_i (p_{ik})^{w_{ij}}, \quad (1)$$

where $\prod_i p_{ik}$ is the product of the probability values derived from the profiles shown in Fig. 6, with the index (k) denoting a specific parent muon. The meaning of the individual parameters q_j , w_{ij} is explained below.

There are two considerations that drive us to building the likelihood in this way. First, each candidate can be associated to a large number of preceding muons, i.e. for each muon-candidate pair, an individual likelihood value \mathcal{L}_k has to be computed. Second, the discrimination parameters feature different levels of indicative force for tagging ^{11}C . Moreover, some are statistically interdependent, some on the other hand contradictory: for instance, while the distances to the neutron vertex and projection, $p(d_n)$ and $p(d_{np})$, are good ^{11}C indicators on their own, their product is likely to be large only for candidates at similar distance from both and thus highly misleading for the purpose of ^{11}C tagging.

To address these issues, the partial probabilities p_i are weighted by exponents w_i that re-scale their relative contribution to the calculation of \mathcal{L} . Note that the introduction of the weights w_i does preserve the shapes of the individual probability profiles (p_i) shown in Fig. 6. In addition, we consider different implementations (j) of the final likelihood, each one putting emphasis on a particular ^{11}C tagging strategy by re-weighting/excluding different sets of discrimination parameters (i):

1. **Proximity to a neutron vertex** is given a special emphasis while the distance to the next projection point is not included.
2. **Proximity to a neutron projection point** puts the tagging priority opposite of (1).
3. **Presence of neutrons in the MW** bears the complication that no vertex reconstruction is available for neu-

Table 1 The case-dependent likelihood implementations ($j = 1..4$) for TFC (see text): probability offsets q_j and weight factors w_{ij} are listed

j	Emphasis	q_j	$w(d_\mu)$	$w(n_n)$	$w(n_{\mu n})$	$w(d_n)$	$w(d_{np})$	$w(\frac{dE}{dx})$	$w(\Delta t)$
1	Closest n vertex	0.025	1	1	0	2	0	0	1
2	Closest n proj.	0.436	1	1	0	0	1	0	1
3	Neutrons in MW	0.740	1.5	0	1	0	0	1	1
4	High n -mult.	0.370	1.5	1	0	0	0	1	1

trons in the MW. Thus, only the presence of neutrons and muon-related information is regarded.

4. Muons featuring high neutron-multiplicities are usually very luminous, again impacting spatial reconstruction for the subsequent neutron vertices. Therefore, vertex information is not regarded for $n_n > 20$, given instead larger emphasis to muon-related parameters.

The corresponding weights w_{ij} are listed in Table 1 and are set to 0 if a specific p_i is not to contribute to the likelihood.

Subsequently, the different implementations (j) are evaluated and the variant providing the highest likelihood value \mathcal{L}_{jk} is selected. To permit the direct comparison of the \mathcal{L}_{jk} , we have defined for each variant (j) an optimum working point for what regards ^{11}C tagging efficiency and ^{11}C -depleted exposure. These figures of merit are introduced in Sect. 3.4. In order to compensate the differences in absolute scale, we apply a posteriori a scaling factor q_j to each likelihood \mathcal{L}_{jk} , permitting a direct comparison between the four implementations (j).

This procedure provides a single likelihood value \mathcal{L}_k for a specific μ -candidate pair. Also the likelihoods obtained for all possible parent muons (k) are compared and again the maximum value is chosen. Thus, the final likelihood \mathcal{L} for a specific ^{11}C candidate is

$$\mathcal{L} = \max_k [\max_j \mathcal{L}_{jk}] \tag{2}$$

All numerical values of the parameters listed in Table 1 have been optimized using the 2012 data set regarding tagging efficiency and preserved exposure (see Sect. 3.4). Like the partial probabilities p_i , the resulting \mathcal{L} thus cannot be considered a likelihood in a strict mathematical sense. However, a high value of \mathcal{L} is indeed a reliable indicator for the cosmogenic origin of an event.

3.2.3 Application as ^{11}C event classifier

Differently from the HC-TFC, the LH-TFC provides the likelihood as a continuous variable for classification as a ^{11}C or *neutrino* event. The distribution of \mathcal{L} as a function of the visible energy is shown in Fig. 7. Indeed, two event categories are discernible: a *neutrino* band centered slightly below $\mathcal{L} \sim 10^{-4}$ and stretching over all visible energies;

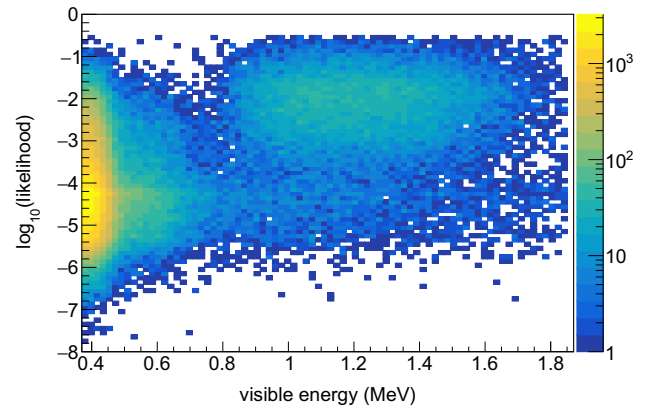


Fig. 7 Distribution of TFC likelihood \mathcal{L} as a function of visible energy. Two event populations can be identified. The upper band centered around $\mathcal{L} \sim 10^{-2}$ are ^{11}C decays concentrated in the ^{11}C energy range. Contrariwise, the lower band centered on $\mathcal{L} \sim 10^{-4}$ consists of *neutrino* events. It is worth to note that the large spectral excess of *neutrino* events at low energies only seems to suggest a vertical structure while it is in fact still centered around the same average value of $\mathcal{L} \sim 10^{-4}$. Based on the 2012 data set

and a smaller cloud of events centered on $\mathcal{L} \sim 10^{-2}$ and concentrated in the energy range 0.75–1.87 MeV that can be associated with ^{11}C events. This is a demonstration that the likelihood \mathcal{L} is indeed a potent classifier for ^{11}C events. Depending on the application, regions in this plot can be selected to either obtain a very pure ^{11}C or *neutrino* sample. The standard application for neutrino analyses is the introduction of an energy-independent cut at $\mathcal{L}_0 = 10^{-4}$ that removes the majority of the ^{11}C events, while keeping about 70% of the *neutrino* events. Analogously to the HC-TFC, this fraction is somewhat reduced when the Full Volume Vetoes discussed in the next section are applied in parallel.

3.3 Full volume regions

There are two situations in which the information on the parent muon and neutron events may be lost: during run breaks and during saturating muon events. These lead both algorithms to apply Full Volume Vetoes (FVV) to effectively reject potential ^{11}C events. We briefly describe them here.

Run breaks. The Borexino data taking is organized in runs with a programmed duration of 6h although they can be shorter if automatically or manually restarted upon prob-

lems. The gap between runs can vary from a few seconds (ordinary restart) to one hour (weekly calibration) or exceptionally to a few days in case of major problems (blackouts, hardware failures). The gap between runs is a time where unrecorded muons could potentially generate ^{11}C in the scintillator which later decays after the run is restarted. TFC algorithms apply a FVV called After-Run-Break (ARB) veto: all events at the beginning of each run are tagged as potential ^{11}C without looking for the parenthood of muon-neutron pairs. The duration of the ARB veto is generally taken to be $\Delta t_{\text{ARB}} = 10 + 60 \cdot (1 - e^{-3\Delta t_{\text{gap}}})$ min, where Δt_{gap} is the duration of the gap to the end of the previous run in minutes.

Showering muons. Muons crossing several meters of scintillator and/or inducing hadronic showers are potentially generating multiple ^{11}C along the track. At the same time the high amount of light can provoke the saturation effects described in Sect. 2. This affects the position reconstruction of subsequent neutrons and can reduce the neutron detection efficiency in very extreme cases. Therefore, both algorithms apply a FVV after events with number of empty boards $N_{\text{EB}} > N_{\text{EB}}^{\text{thr}}$ and $n_n > 0$ for 4–5 $\tau_{^{11}\text{C}}$. We observe that the average value of N_{EB} for non saturating events is slowly changing through the data set, mostly increasing due to the diminishing number of live channels. In addition a few rearrangement of the electronics and cabling also impacted on this parameter. In order to take this into account we adjust $N_{\text{EB}}^{\text{thr}}$ to a value above the bulk of the distribution in each subset of data defined by the same electronics configuration and let the outliers trigger the FVV.

It should be noted that the data within FVV are considered when TFC algorithms search for muon-neutron pairs.

FVV are very important to keep the tag efficiency high but have a high price in terms of exposure. This consideration led to the development of the alternative BI algorithm described in Sect. 4.

3.4 TFC comparison

We present the performance of the two TFC algorithms described in terms of two desiderata. The TFC technique aims to maximize the ^{11}C tagging efficiency in the ^{11}C -enriched data set, while retaining the maximum exposure fraction in the ^{11}C -depleted set. The trade-off between those two competing requirements is ultimately decided by several parameters for the HC-TFC and by a single threshold \mathcal{L}_0 in case of the LH-TFC. While the working point we present here is the result of a careful optimization, the best parameter combinations for each neutrino analysis can be determined independently and may be to some extent different than what is presented here.

Exposure fraction. The first figure of merit is the fraction of the exposure $\varepsilon_{\text{depleted}}$ that remains in the ^{11}C -depleted data set

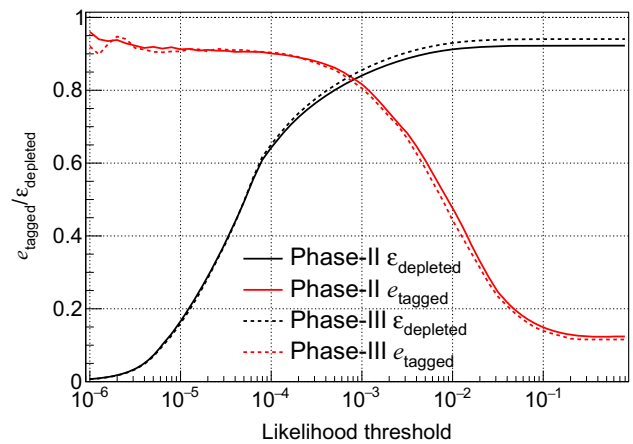


Fig. 8 Performance of the LH-TFC as a function of its threshold value \mathcal{L}_0 . ^{11}C tagging efficiency (e_{tagged}) and the exposure of the ^{11}C -depleted data set ($\varepsilon_{\text{depleted}}$) can be adjusted to the need of a particular analysis. For solar analysis, the default choice is $\mathcal{L}_0 = 10^{-4}$

after TFC application. As laid out in Sect. 3, $\varepsilon_{\text{depleted}}$ is determined statistically by inserting random vertices (both position and time) into the experimental data stream and counting the fraction that is selected by the TFC criteria for inclusion in the ^{11}C -depleted data set.

Tagging efficiency. The second figure of merit is the tagging efficiency e_{tagged} . Its value is estimated based on a relatively narrow energy range from 1.2 to 1.4 MeV that is selected to minimize the presence of all spectral components but ^{11}C . In order to correct for the loss of live exposure in the depleted data set, e_{tagged} is defined as:

$$e_{\text{tagged}} = 1 - \frac{{}^{11}\text{C}_{\text{untagged}}}{{}^{11}\text{C}_{\text{total}}} \cdot \frac{1}{\varepsilon_{\text{depleted}}} \tag{3}$$

where ${}^{11}\text{C}_{\text{total}}$ and ${}^{11}\text{C}_{\text{untagged}}$ are the total number of events and that of untagged events, respectively. It is worth noting that e_{tagged} is actually a lower limit, as spectral analysis suggest a presence of $\sim 5\%$ of non- ^{11}C events in the evaluated energy range.

TFC tuning. For both TFC algorithms a large flexibility of performance can be obtained tuning the selection parameters. In the case of the LH-TFC this is particularly easy as tagging efficiency and exposure can be varied as a function of the likelihood threshold \mathcal{L}_0 as shown in Fig. 8. While a value of $\mathcal{L}_0 = 10^{-4}$ provides high tagging power for a moderate exposure loss that is ideal for solar neutrino analyses, a choice of $\mathcal{L}_0 = 10^{-2}$ would mean that a much larger fraction of the exposure would be preserved while only half of the ^{11}C candidates are tagged. In this case, the corresponding ^{11}C -enriched sample would almost exclusively consist of real ^{11}C decay events.

Table 2 Tagging efficiency (calculated with Eq. (3)) and exposure fraction for the two TFC approaches. For the LH-TFC, \mathcal{L}_0 is set at 10^{-4} . The statistical uncertainty on all values of e_{tagged} is 0.5%, while it is negligible on $\epsilon_{\text{depleted}}$ values

	Phase-II (%)	Phase-III (%)
Hard cut		
Tagging efficiency (e_{tagged})	90.2	90.7
Exposure fraction ($\epsilon_{\text{depleted}}$)	63.3	63.6
Likelihood		
Tagging efficiency (e_{tagged})	89.8	90.1
Exposure fraction ($\epsilon_{\text{depleted}}$)	64.7	65.6

TFC performance. The performance results are presented in Table 2 and are based on the data of Borexino Phase-II and Phase-III introduced in Sect. 2. Although the performance is essentially independent from the choice of the fiducial volume, for consistency we use the same fiducial volume of the *pep* and CNO neutrino analyses (Sect. 2).

Both algorithms return a tagging efficiency of about 90% while providing a ^{11}C -depleted exposure fraction of $\sim 64\%$. While HC-TFC provides slightly better tagging, LH-TFC saves somewhat more exposure. In recent solar neutrino analyses [6], both algorithms have been applied separately to investigate potential differences in the corresponding spectral fit results and to assess systematic uncertainties.

Time stability. We also investigated the stability of the TFC performance on a yearly bases throughout Borexino Phase-II and III and extending till the end of 2020. In Figs. 9 and 10 the ^{11}C tagging efficiency and the exposure fraction are shown for both TFC implementations. We attribute the slightly better performance of both TFC methods during Phase-III to a more continuous data taking which implies a lower impact of ARB cuts. Moreover, we observe a drift in LH-TFC performance during Phase-III, trading slightly lower tagging efficiencies for higher exposure fractions compared to HC-TFC. In principle, this trend could be compensated by a year-by-year adjustment of the likelihood threshold \mathcal{L}_0 .

4 ^{11}C burst identification

The implementations of the TFC discussed so far have been extensively tested and validated and were part of several of the Borexino solar- ν analyses.

Here, we describe a novel method that further improves the identification of ^{11}C background events in Borexino and is potentially applicable as well to other large volume liquid scintillator detectors.

Most of ^{11}C is produced not by primary muon spallation but by secondary processes like (p, n) reactions occurring in muon-induced hadronic showers. As a consequence, multiple

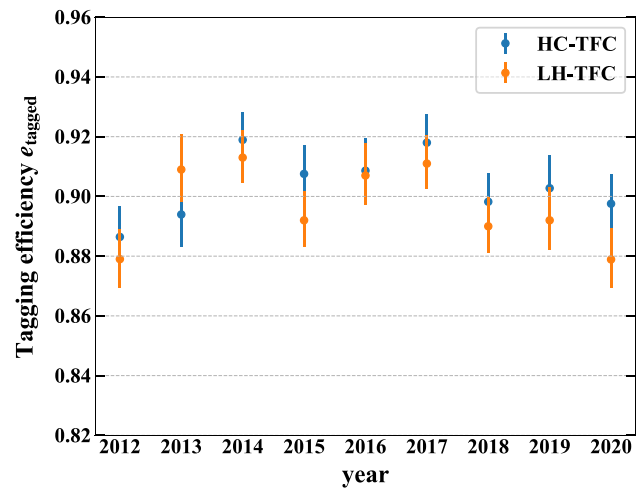


Fig. 9 Tagging efficiency for HC-TFC (blue dots) and LH-TFC (orange dots) year by year

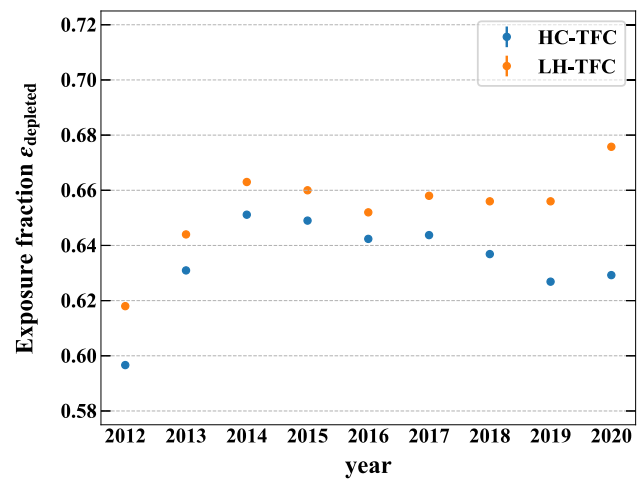
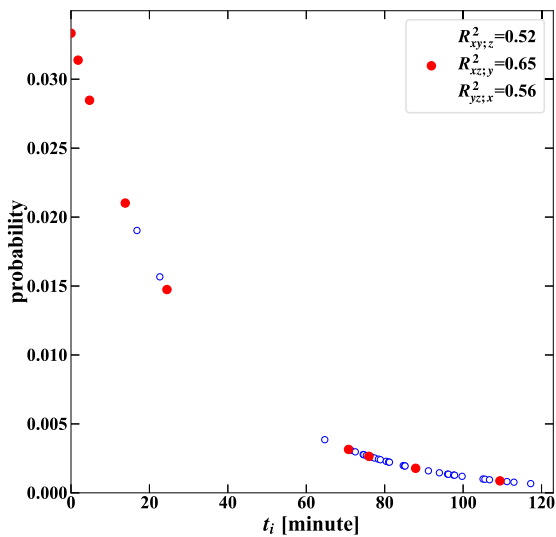


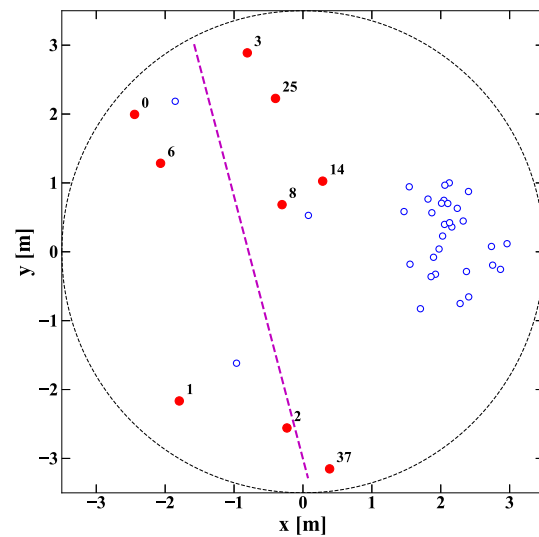
Fig. 10 Exposure fraction of the depleted sample for HC-TFC (blue dots) and LH-TFC (orange dots) year by year

^{11}C nuclei are produced simultaneously in the detector and subsequently decay over a couple of hours. Consequently, identification of a *burst* of possible ^{11}C candidates correlated in space (i.e. along a hypothetical muon track) and time (several ^{11}C lifetimes) can provide a tag independent of any information on parent muon and neutron events. This is of particular interest since this Burst Identification (BI) tag will be available even if the initiating muon and/or subsequent neutron captures have been missed or misidentified by the data acquisition. Therefore, when applied in combination with either of the two TFC algorithms, the BI tag offers a more convenient alternative to the Full Volume Veto (FVV) introduced in Sect. 3.3.

The BI algorithm presented here proceeds first to identify groups of ^{11}C candidates, i.e. events in the correct energy window, with sufficient space and time correlation with each



(a) Probability distribution from ^{11}C decay time in a 4τ time window.



(b) Spatial distribution in the xy -plane.

Fig. 11 Sample group of events considered for building a *burst*. Red full dots highlight the events that are finally selected as spatially correlated, with the multi-correlation factors R^2 reported in the legend in **a**. The numbers in **b** reflect the time ordering (for clarity only some are shown). The dashed circle shows the radius of 3.5 m used to pre-

select *burst* candidates. The dashed magenta line is the projection of the regression line. Open blue dots represent events not correlated with the first *burst*. Note that both panels show a late and spatially dense population belonging most likely to a second coinciding ^{11}C *burst* that is distinguished as separate by the algorithm

other. Once a *burst* has been identified, it is used to tag all events that show a correlation with the events that defined it.

4.1 Step 1: finding the *burst*

Time correlation. The algorithm starts considering successively all ^{11}C candidates searching for time correlations. Candidate events are defined by a visible energy within 0.75–1.87 MeV corresponding to the ^{11}C spectrum. At this stage a very loose fiducial volume cut is applied: all events within a 3.5 m radius are considered. This has been optimized for excluding most of the external background events, without losing many ^{11}C events that could be helpful in finding a *burst*. For each candidate, a subsequent time window corresponding to $4\tau_{11\text{C}}$ (~ 2 hours) is opened. In order to form a potential *burst*, at least four ^{11}C candidates have to be present inside the window (including the initial one). Figure 11a shows a sample of events considered in the definition of a *burst* to whom we have associated a probability based on the exponential decay time of ^{11}C . We build a cumulative probability summing these values for all candidates, $\sum_i 1/\tau_{11\text{C}} \cdot \exp(-(t_i - t_0)/\tau_{11\text{C}})$, that is required to exceed a given threshold p_{th} . In the case of a run interruption within the observation window, this threshold is adjusted to take into account the dead time.

In some rare cases, it may happen that a *burst* falls within the selection window of a preceding unrelated ^{11}C candidate.

The latter can then be mistaken as the first event and seriously distort the construction of the *burst*. In order to prevent this, we check whether the time distribution of events inside the selection window follows an exponential decay law. Specifically, if the time gap between first and second candidate is disproportionately large, the first event is considered unrelated to the *burst* and discarded.

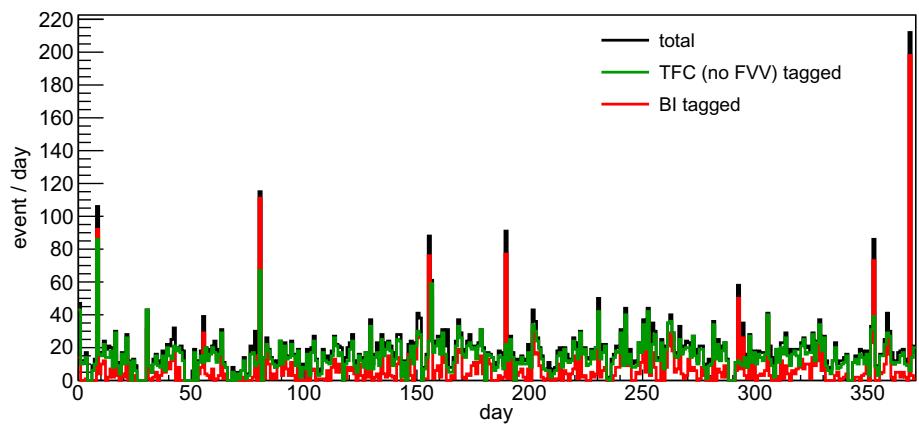
Spatial correlation. After a time correlation is established, the algorithm proceeds to check if any spatial correlation exists among the time-selected ^{11}C candidates. The basic assumption is that ^{11}C is produced and decays along a straight line (i.e. a muon track). Events that are far from this line are likely to be accidental time coincidences and must not be included in the *burst*. The hypothesis is tested via the correlation coefficients of the events coordinates upon projecting on the coordinate planes.

Figure 11b shows the projection on the (xy) -plane of the sample events. Since the line is a freely oriented three-dimensional object, the use of classic Pearson correlation coefficients $\rho_{ij} = \text{cov}(i, j)/\sigma_i\sigma_j$ is not sufficient. Instead, the correlation is checked based on the three multi-correlation factors

$$R^2_{ij;k} = (\rho^2_{ik} + \rho^2_{jk} - 2\rho_{ij}\rho_{ik}\rho_{jk})/(1 - \rho^2_{ij}), \quad (4)$$

with (i, j, k) denoting the different combinations of spatial coordinates (x, y, z) that correspond to the coordinate planes.

Fig. 12 Events per day in the ^{11}C region inside the fiducial volume, in a year-long window of Phase-III. Black: total daily rate. Green: rate of events tagged by the TFC with the Full Volume Veto (FVV) switched off (see text for details). Red: rate of events tagged by the Burst Identification (BI) algorithm. Spikes in the rate correspond to *bursts* of ^{11}C



Starting values for the $R_{ij;k}^2$ are computed based on the first four time-selected events (denoted as 0–3 in Fig. 11). For each following event, the $R_{ij;k}^2$ values are re-computed after adding the event to the *burst*. The event is considered aligned with the track if: (1) all new $R_{ij;k}^2$ values remain larger than a given lower limit R_{lim}^2 ; (2) the two largest new $R_{ij;k}^2$ values do not differ from the previous values by more than a tolerance margin R_{tol}^2 . Upon meeting both conditions the event is accepted to belong to the *burst* and the new $R_{ij;k}^2$ values become the new reference. In Fig. 11 the events finally included in the *burst* after spatial selection are indicated by red filled circles. The final correlation factors $R_{\text{burst};k}^2$ provide the thresholds for the application of the tagging (Sect. 4.2). The procedure is repeated upon sorting the events after their distance to the initial candidate and the case showing a stronger correlation is preferred.

For the performance evaluation presented in Sect. 4.3, the values of the parameters discussed above were chosen to be $p_{\text{th}} = 0.12$, $R_{\text{lim}}^2 = 0.1$, and $R_{\text{tol}}^2 = 40\%$ (20%) for Phase-II (Phase-III). Although these parameter values are the result of a careful calibration of the algorithm, a different choice could be made in the context of a specific neutrino analysis in order to achieve a different trade-off between tagging efficiency and exposure fraction.

4.2 Step 2: tagging the ^{11}C

Once a *burst* is identified by the events selected as in Sect. 4.1, all events within the *burst* time window, regardless of their energy, are checked for their affiliation with the *burst* and tagged as ^{11}C .

The track underlying the *burst* is found by a Theil–Sen regression [27,28] returning the direction and the impact parameter based on the median of the tracks passing through all possible pairs of events of the *burst* (this method is less influenced by outliers compared to those based on the mean). In Fig. 11b, the projection of the track on the (x,y) -plane is shown as a magenta dashed line.

Thereafter, events within the *burst* time window are ordered by their distance to the track. Starting with the four closest events and then subsequently adding further events one by one, the algorithm proceeds to calculate the multi-correlation factors as in Eq. (4). Newly added events are considered correlated with the track as long as all three $R_{ij;k}^2$ remain above the corresponding $R_{\text{burst};k}^2$. Events may be tested for correlation multiple times if they fall within the time window of more than one track. An event is tagged as ^{11}C by the BI algorithm if it is correlated within at least one track.

Similarly to the TFC, the BI algorithm calculates the exposure fractions via toy Monte Carlo. Simulated events are uniformly distributed inside the fiducial volume and their correlation with all *burst* tracks is computed. The fraction of events for which $R_{ij;k}^2 < R_{\text{burst};k}^2$ for all tracks corresponds to the live exposure fraction.

4.3 BI performance

As discussed in Sect. 3.3, the TFC algorithms can only identify ^{11}C decays for which the parent muon and subsequent neutrons have been detected. There are situations following breaks in data acquisition and especially very bright muon events saturating the electronics where this requirement is not satisfied. The Full Volume Vetoes (FVV) that have to be introduced to compensate for this deficiency are very costly in exposure for the ^{11}C -depleted sample. Here we demonstrate that the same task is more efficiently fulfilled by the BI algorithm.

Figure 12 shows the daily rate of events in the ^{11}C energy range and inside the fiducial volume (Sect. 2) over the course of a sample period of 1 year. It also shows the rates of those events identified as ^{11}C by the HC-TFC without FVV and of those identified by the BI algorithm. While this basic TFC identifies most of the ^{11}C events during regular data taking, it misses a larger fraction of events during the bursts due to exceptional high-multiplicity ^{11}C production. This behaviour is expected since ^{11}C -bursts are accompanied by severe satu-

Table 3 Performance of the HC-TFC without Full Volume Veto (FVV) before and after the combination with the Burst Identification (BI) tag, compared with the standard HC-TFC already shown in Table 2. The statistical uncertainty on all values of ϵ_{tagged} is 0.5%, while it is negligible on $\epsilon_{\text{depleted}}$ values

HC-TFC		Phase-II (%)	Phase-III (%)
no FVV	Tagging efficiency (ϵ_{tagged})	79.5	77.6
	Exposure fraction ($\epsilon_{\text{depleted}}$)	70.8	69.7
no FVV + BI	Tagging efficiency (ϵ_{tagged})	89.3	90.4
	Exposure fraction ($\epsilon_{\text{depleted}}$)	68.3	66.7
with FVV (standard)	Tagging efficiency (ϵ_{tagged})	90.2	90.7
	Exposure fraction ($\epsilon_{\text{depleted}}$)	63.3	63.6

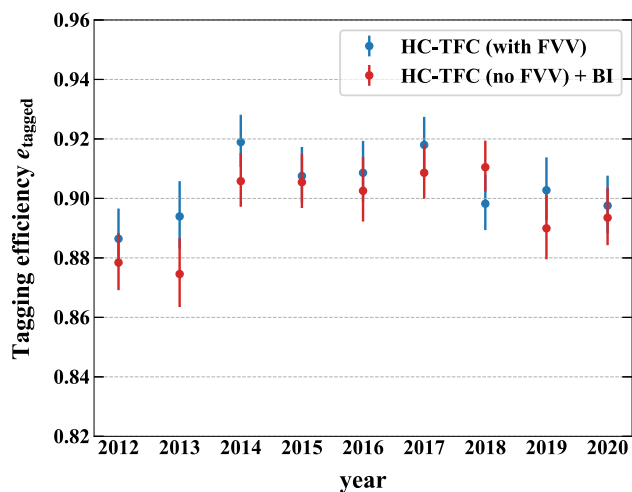


Fig. 13 Tag efficiency year by year for TFC (with Full Volume Veto, FVV) and TFC (no FVV) + Burst Identification (BI), respectively in blue and red

ration of the electronics. On the other hand, the BI algorithm shows heightened efficiency for these peak regions.

Hence, an event is assorted to the ^{11}C -enriched sample in case it is tagged either by the TFC (without FVV)³ or by the BI algorithm. Similarly, we evaluate the exposure based on a shared toy Monte Carlo generation. Results for the combined tagging efficiency and the ^{11}C -depleted exposure fraction are shown in Table 3. The combination shows a tagging efficiency similar to the full TFC (i.e. including FVV), but as expected it preserves a $\sim 3\text{--}5\%$ larger exposure for the ^{11}C -depleted data set.

Similarly to Sect. 3.4, we have studied the time stability of the combined tags. The results are shown in Figs. 13 and 14 for Phase-II, Phase-III, and extending until the end of 2020.

5 Conclusions

Low-energy neutrino detection in organic liquid scintillators suffers from spallation of carbon nuclei by cosmic muons.

³ We chose HC-TFC as reference in this context, but the combination of LH-TFC and BI performs similarly.

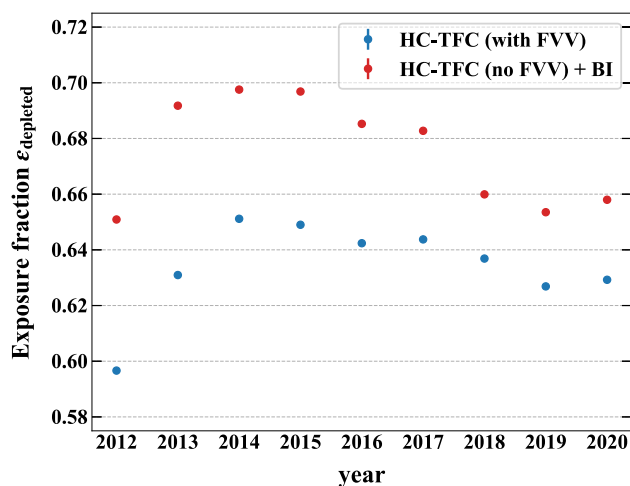


Fig. 14 Exposure fraction of the depleted sample year by year for TFC (with FVV) and TFC (no FVV) + BI, respectively in blue and red

We have presented the state-of-the-art tagging and veto techniques for the most frequent spallation product ^{11}C that have been developed in Borexino in the context of solar *pep* and CNO neutrino analyses.

We have described the Three-Fold Coincidence (TFC) technique that makes use of time and spatial correlation among parent muons, neutrons, and ^{11}C decay events. We have given details of its two current implementations, a Hard-Cut and a Likelihood-based approach. Scanning the data of Borexino Phases II (2012–2016) and III (2016–2020), we showed that both methods return stable results over a long time period. Typical tagging efficiencies exceed 90%, while the ^{11}C -depleted data set retains about 65% of the exposure.

Finally, we discussed the ^{11}C -burst tagging technique that exploits the fact that ^{11}C is often created in bursts by major spallation events. This novel approach searches for clusters of ^{11}C candidates that feature an exponentially decaying time behavior and are closely aligned along a track. We showed how this technique can be combined with the TFC and determined the corresponding improvement in surviving exposure fraction (66.7% vs. 63.6%) for essentially the same tagging power on Phase-III data. Given the net benefit, the method is intended for use in future solar neutrino analyses of Borexino.

While native to Borexino, these methods can be easily adjusted to the identification and veto of other isotopes or to alternative experimental setups, and thus may be of benefit for a wide range of present and future low-energy neutrino and other low-background experiments.

Acknowledgements The Borexino program is made possible by funding from Istituto Nazionale di Fisica Nucleare (INFN) (Italy), National Science Foundation (NSF) (USA), Deutsche Forschungsgemeinschaft (DFG) and Helmholtz-Gemeinschaft (HGF) (Germany), Russian Foundation for Basic Research RFBR (Grant 19-02-00097 A), RSF (Grant 21-12-00063) (Russia), and Narodowe Centrum Nauki (NCN) (Grant no. UMO 2017/26/M/ST2/00915) (Poland). This research was supported in part by PLGrid Infrastructure. We acknowledge the generous hospitality and support of the Laboratory Nazionali del Gran Sasso (Italy).

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors' comment: Any request for data can be addressed to spokesperson-borex@lngs.infn.it.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. Funded by SCOAP³.

References

- G. Bellini et al., Phys. Rev. Lett. **107**, 141302 (2011). <https://doi.org/10.1103/PhysRevLett.107.141302>
- G. Bellini et al., Phys. Rev. D **82**, 033006 (2010). <https://doi.org/10.1103/PhysRevD.82.033006>
- G. Bellini et al., Phys. Rev. Lett. **108**, 051302 (2012). <https://doi.org/10.1103/PhysRevLett.108.051302>
- G. Bellini et al., Nature **512**(7515), 383 (2014). <https://doi.org/10.1038/nature13702>
- M. Agostini et al., Nature **562**(7728), 505 (2018). <https://doi.org/10.1038/s41586-018-0624-y>
- M. Agostini et al., Phys. Rev. D **100**(8), 082004 (2019). <https://doi.org/10.1103/PhysRevD.100.082004>
- M. Agostini et al., Phys. Rev. D **101**(6), 062001 (2020). <https://doi.org/10.1103/PhysRevD.101.062001>
- M. Agostini et al., Nature **587**, 577 (2020). <https://doi.org/10.1038/s41586-020-2934-0>
- M. Agostini et al., Eur. Phys. J. C **80**(11), 1091 (2020). <https://doi.org/10.1140/epjc/s10052-020-08534-2>
- M. Deutsch. Proposal to NSF for a Borexino Muon Veto. Massachusetts Institute of Technology (1996)
- C. Galbiati et al., Phys. Rev. C **71**, 055805 (2005). <https://doi.org/10.1103/PhysRevC.71.055805>
- G. Bellini et al., Phys. Rev. D **89**(11), 112007 (2014). <https://doi.org/10.1103/PhysRevD.89.112007>
- S. Abe et al., Phys. Rev. C **81**, 025807 (2010). <https://doi.org/10.1103/PhysRevC.81.025807>
- V. Lozza, *The SNO+ Experiment* (2019), p. 313–328. https://doi.org/10.1142/9789811204296_0019
- F. An et al., J. Phys. G **43**(3), 030401 (2016). <https://doi.org/10.1088/0954-3899/43/3/030401>
- J.F. Beacom, M.R. Vagins, Phys. Rev. Lett. **93**, 171101 (2004). <https://doi.org/10.1103/PhysRevLett.93.171101>
- K. Abe et al., (2011)
- S.W. Li, J.F. Beacom, Phys. Rev. D **92**(10), 105033 (2015). <https://doi.org/10.1103/PhysRevD.92.105033>
- M. Askins et al., Eur. Phys. J. C **80**(5), 416 (2020). <https://doi.org/10.1140/epjc/s10052-020-7977-8>
- G. Alimonti et al., Nucl. Instrum. Methods A **600**, 568 (2009). <https://doi.org/10.1016/j.nima.2008.11.076>
- G. Bellini et al., JCAP **08**, 049 (2013). <https://doi.org/10.1088/1475-7516/2013/08/049>
- G. Bellini et al., JINST **6**, P05005 (2011). <https://doi.org/10.1088/1748-0221/6/05/P05005>
- M. Agostini et al., JCAP **02**, 046 (2019). <https://doi.org/10.1088/1475-7516/2019/02/046>
- P.A. Sanchez et al., Phys. Lett. B **716**, 401 (2012). <https://doi.org/10.1016/j.physletb.2012.08.052>
- Y. Abe, et al., JHEP **10**, 086 (2014). [https://doi.org/10.1007/JHEP02\(2015\)074](https://doi.org/10.1007/JHEP02(2015)074) (Erratum: JHEP **02**, 074 (2015))
- H. de Kerret et al., JHEP **11**, 053 (2018). [https://doi.org/10.1007/JHEP11\(2018\)053](https://doi.org/10.1007/JHEP11(2018)053)
- H. Theil, Nederl. Akad. Wetensch. **53**, 386 (1950)
- P.K. Sen, J. Am. Stat. Assoc. **63**(324), 1379 (1968). <https://doi.org/10.2307/2285891>