

Probing cosmological particle production and pairwise hotspots with deep neural networks

Taegyun Kim,^{1,*} Jeong Han Kim,^{2,†} Soubhik Kumar,^{3,4,‡} Adam Martin,^{1,§} Moritz Münchmeyer,^{5,||} and Yuhsin Tsai^{1,¶}

¹*Department of Physics, University of Notre Dame, Indiana 46556, USA*

²*Department of Physics, Chungbuk National University, Cheongju, Chungbuk 28644, Korea*

³*Theoretical Physics Group, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

⁴*Berkeley Center for Theoretical Physics, Department of Physics, University of California, Berkeley, California 94720, USA*

⁵*Department of Physics, University of Wisconsin-Madison, 1150 University Avenue, Madison, Wisconsin 53706, USA*



(Received 14 April 2023; accepted 2 August 2023; published 22 August 2023)

Particles with masses much larger than the inflationary Hubble scale, H_I , can be pair-produced nonadiabatically during inflation. Due to their large masses, the produced particles modify the curvature perturbation around their locations. These localized perturbations eventually give rise to localized signatures on the cosmic microwave background (CMB), in particular, pairwise hotspots (PHSs). In this work, we show that convolutional neural networks (CNNs) provide a powerful tool for identifying PHSs on the CMB. While for a given hotspot profile a traditional matched filter analysis is known to be optimal, a neural network learns to effectively detect the large variety of shapes that can arise in realistic models of particle production. Considering an idealized situation where the dominant background to the PHS signal comes from the standard CMB fluctuations, we show that a CNN can isolate the PHSs with $\mathcal{O}(10)\%$ efficiency even if the hotspot temperature is $\mathcal{O}(10)$ times smaller than the average CMB fluctuations. Overall, the CNN search is sensitive to heavy particle masses $M_0/H_I = \mathcal{O}(200)$, and constitutes one of the unique probes of very high energy particle physics.

DOI: [10.1103/PhysRevD.108.043525](https://doi.org/10.1103/PhysRevD.108.043525)

I. INTRODUCTION

An era of cosmic inflation [1–3] in the primordial Universe remains an attractive paradigm to explain the origin of (approximately) scale invariant, Gaussian, and adiabatic primordial perturbations, inferred through cosmic microwave background (CMB) and large scale structure observations. This inflationary era can be characterized by a rapid expansion of spacetime, controlled by an approximately constant Hubble scale H_I . Excitingly, based on the current constraints, H_I can be as large as 5×10^{13} GeV [4]. This fact, coupled with the feature that particles with masses up to order H_I can get quantum mechanically

produced during inflation, makes the inflationary era a natural and unique arena to directly probe very high energy particle physics.

There are several classes of mechanisms through which heavy particles, which we label as χ , can be produced during inflation. When their mass $m_\chi \lesssim H_I$, quantum fluctuations of the inflationary spacetime itself can efficiently produce the χ particles. However, for $m_\chi \gg H_I$, this production gets suppressed exponentially as $e^{-\pi m_\chi/H_I}$ [5], and other mechanisms are necessary for efficient particle production to occur.

To illustrate this, we consider the standard slow-roll inflationary paradigm containing an inflaton field ϕ whose homogeneous component we denote by $\phi_0(t)$. Normalization of the primordial scalar power spectrum requires the “kinetic energy” of this homogeneous component to be $|d\phi_0/dt|^{1/2} \approx 60H_I$ [4]. Therefore, heavy particles, if appropriately coupled to the inflaton kinetic term, can be efficiently produced for $m_\chi \lesssim 60H_I$. One class of this involves a coupling of the type $\partial_\mu \phi J^\mu$, where J^μ is a charged current made up of the χ field. For some recent work implementing this idea see, e.g., Refs. [6–14]. In these constructions, heavy particle production happens continuously in time, in a scale-invariant fashion. In other words, the coupling of the

*kim12@nd.edu

†jeonghan.kim@cnu.ac.kr

‡soubhik@berkeley.edu

§amarti41@nd.edu

||muenchmeyer@wisc.edu

¶ytsai3@nd.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by SCOAP³.

inflaton to χ particles does not break the shift symmetry, $\phi \rightarrow \phi + \text{constant}$, of the inflaton.

A different class of mechanisms can lead to particle production at specific times during the inflationary evolution. This can happen if the shift symmetry of the inflaton is broken in a controlled manner, e.g., to a discrete shift symmetry. This breaking of shift symmetry translates into a violation of scale invariance and selects out specific time instant(s) when particle production can occur. Examples of such mechanisms appear in Refs. [15–21], and see Refs. [22,23] for reviews.

A particularly interesting example of this latter mechanism arises in the context of ultraheavy particles with time-dependent masses. More specifically, suppose m_χ varies as a function of ϕ in a way such that, as ϕ passes through a specific point ϕ_* on the inflaton potential at time t_* , $m_\chi(\phi)$ passes through a local minimum. In this case, nonadiabatic χ particle production can occur at time t_* . Following their production, χ particles can again become heavy, $m_\chi \gg |d\phi_0/dt|^{1/2}$, and owing to this large mass they can backreact on the inflationary spacetime, contributing to the curvature perturbation around their locations.

We can describe the effects of these additional curvature perturbations qualitatively in the following way, leaving the details for the next section. Following their production, the perturbations exit the horizon when their wavelengths become larger than $1/H_I$ and become frozen in time. After the end of inflation, they eventually reenter the horizon and source additional under- or overdensities in the thermal plasma in the radiation dominated Universe. Overdense regions, for example, would trap more plasma, and therefore would emit more photons at the time of CMB decoupling.¹ Therefore, we would observe localized regions on the sky where CMB would appear hotter than usual. As we will discuss below, the sizes of these localized “spots” are determined by the size of the comoving horizon, η_* , at the time of particle production t_* . While η_* can take any value, for concreteness we will consider $\eta_* \sim 100$ Mpc in this work. This implies that the localized spots would subtend $\sim 1^\circ$ on the CMB sky.

The next question one may ask is what is an efficient strategy to look for such signatures. Since this scenario is associated with a violation of scale invariance, characterized by η_* , one would expect to see “features” on the CMB power spectrum or even higher-point correlation functions. However, in the regime we focus on, the total number of produced χ particles is still small to the extent that the CMB power spectrum is minimally affected, as we explicitly check later. On the other hand, the spots can still be individually bright enough such that we can look for them

directly in position space. Higher N -point functions can be more sensitive than the power spectrum, but recent work [24] showed that an analysis of higher N -point functions due to particle production can be resummed to the much simpler profile search in position space, assuming the density of profiles is low enough so that they do not usually overlap. In this case a profile search will always be statistically optimal. This class of signatures in the context of heavy particle production were discussed in Refs. [25,26], and in Ref. [27] the associated CMB phenomenology was described and a simple “cut-and-count” search strategy was developed. Using the cut-and-count strategy, Ref. [27] constrained the parameter space of ultraheavy scalars and illustrated regions where a position space search is more powerful than power spectrum-based searches.

In more detail, Ref. [27] considered a single instance of particle production during the time when CMB-observable modes exit the horizon. Conservation of momentum implies that such heavy particles are produced in pairs. However, owing to their large mass, the particles do not drift significantly following their production, and it was argued that the separation between the two particles forming a pair can be taken to be a uniformly random number between 0 and η_* . Finally, it was shown that the coupling g of χ to the inflaton determines how hot/cold the associated spot on the CMB is with the heavy particle mass m_χ determining the total number of such spots on the sky. To summarize, the three parameters determining the hot/cold spot phenomenology are $\{g, m_\chi, \eta_*\}$, as will be reviewed in more detail in the next section. While both cold or hot spots can arise depending on the value of η_* , for the choices of η_* in this work, only hotspots will appear on the CMB. Therefore, we will often be referring to these localized spots as hotspots, in particular as pairwise hotspots (PHSs) since the spots appear in pairs.

In the present work, we improve upon Ref. [27] in several important ways. First, in Ref. [27] we only considered hotspots that lie within the last scattering surface, with a thickness of $\Delta\eta \approx 19$ Mpc [28]. In this work we adopt a more realistic setup and include hotspots that are distributed in a larger region around the last scattering surface. We take this region to have a thickness of $2\eta_*$ and we show in Sec. II how hotspots lying outside the $\Delta\eta$ shell can still affect the CMB. The overall signature of PHSs then changes nontrivially. For instance, with the improved treatment we can have one spot of a pair lying on the CMB surface, while the other can lie off the CMB surface, leading to an asymmetric signal.

Second, we develop a neural-network-based search for the hotspot profiles. In principle, a neural network is not necessary to search for a profile of known shape which is linearly added to the Gaussian background. In this case, the standard method of constructing a so-called matched filter can be shown to be the optimal statistic to detect the profile (see, e.g., [24]). Matched filter-based searches for radially

¹To be more accurate, one also needs to take into account the gravitational redshift of the photons as they climb out of the gravitational potential wells. We will compute this effect in the next section.

symmetric profiles in the CMB have been previously reported for example in [29–31], with the physical motivation of searching for inflationary bubble collisions. Various matched filters have also been used in the Planck anisotropy and statistics analysis [32,33] without finding a significant excess. However, the signal that we are looking for here is more complicated. Profiles come in pairs (breaking radial symmetry of the profile), they can be overlapping, and, depending on their production time and orientation with respect to the surface of last scattering, their appearance on the CMB changes. While it is in principle possible to cover the entire space of profiles with a very large bank of matched filters, this would be a complicated and computationally challenging approach. A neural network, on the other hand, can learn an effective representation of these filters that interpolates well between all profile shapes, including overlapping ones. We also implement the matched filter method below, and show that in the simplified case with a single profile type, our neural network performs similar to the optimal matched filter.

This work is organized as follows. We first describe a simple model of χ particle production in Sec. II and summarize how the total number of produced particles depends on the model parameters along with various properties of the PHSs. We improve the calculation of the hotspot profiles by taking into account the line-of-sight distance to the location of the hotspots which can be off the CMB surface. In Sec. III, we describe the simulation of the PHS signals and the CMB maps in angular space, assuming that the dominant background to the PHS signal comes from the standard CMB fluctuations. In Sec. IV, we describe the convolutional neural network (CNN) analysis and estimate the sensitivity the CNN can achieve for a PHS search. We then translate this sensitivity to the mass-coupling parameter space of the heavy particles. We also compare the CNN analysis with a matched filter analysis for simplified hotspot configurations. We conclude in Sec. V.

II. PAIRWISE HOTSPOT SIGNALS

To model heavy particle production, we consider a scenario where the mass of χ is inflaton dependent, $m_\chi(\phi)$. Therefore as ϕ moves along its potential, efficient, nonadiabatic particle production can occur if $m_\chi(\phi)$ varies with ϕ rapidly. With a mass term $m_\chi(\phi)^2\chi^2$, pairs of χ particles would be produced, as required by three-momenta conservation. The phenomenology of such heavy particles depend on their mass, coupling to the inflaton, and the horizon size at the time of their production. We now review these properties more qualitatively, referring to Ref. [27] for a more complete discussion.

A. Inflationary particle production

We parametrize the inflationary spacetime metric as

$$ds^2 = -dt^2 + a^2(t)d\vec{x}^2, \quad (1)$$

with the scale factor $a(t) = e^{H_I t}$ and H_I the Hubble scale during inflation that we take to be (approximately) constant. To model particle production in a simple way, we assume $m_\chi(\phi)$ passes through a minimum as ϕ crosses a field value ϕ_* . Then we can expand $m_\chi(\phi)$ near ϕ_* as

$$m_\chi(\phi) = m_\chi(\phi_*) + \frac{1}{2}m_\chi''(\phi_*)(\phi - \phi_*)^2 + \dots, \quad (2)$$

where primes denote derivatives with respect to ϕ . Thus the mass term would appear in the potential as

$$m_\chi(\phi)^2\chi^2 = m_\chi(\phi_*)^2\chi^2 + m_\chi(\phi_*)m_\chi''(\phi_*)(\phi - \phi_*)^2\chi^2 + \dots \quad (3)$$

While away from ϕ_* , $m_\chi(\phi)$ can vary in different ways, most of the important features of particle production are determined by the behavior of $m_\chi(\phi)$ around ϕ_* . For example, the number density of χ particles is determined by $m_\chi(\phi_*)$, as we will see below. Similarly, the spatial profiles of the hotspots on the CMB is determined by the dependence $(\phi - \phi_*)^2 \sim \dot{\phi}_0^2(t - t_*)^2 \sim (\dot{\phi}_0/H_I)^2 \log(\eta/\eta_*)^2$, where we have used the relation between t and conformal time η , $\eta = (-1/H_I)e^{-H_I t}$ (an overdot here denotes a derivative with respect to time). Given the importance of the physics around ϕ_* , we will denote, $m_\chi(\phi_*)^2 \equiv M_0^2$, $m_\chi(\phi_*)m_\chi''(\phi_*) \equiv g^2$, and $\phi_* \equiv \mu/g$, to describe particle production. Thus we will write the Lagrangian for χ as

$$\mathcal{L}_\chi = -\frac{1}{2}(\partial_\mu\chi)^2 - \frac{1}{2}((g\phi - \mu)^2 + M_0^2)\chi^2. \quad (4)$$

As ϕ nears the field value ϕ_* , the mass of the χ field changes nonadiabatically and particle production can occur.

The efficiency of particle production depends on the parameters g , M_0 , and η_* , the size of the comoving horizon at the time of particle production. This can be computed using the standard Bogoliubov approach, and resulting probability of particle production is given by [20,34]

$$|\beta|^2 = \exp\left(-\frac{\pi(M_0^2 - 2H_I^2 + k^2\eta_*^2H_I^2)}{g|\dot{\phi}_0|}\right). \quad (5)$$

The normalization of the scalar primordial power spectrum, in the context of single-field slow-roll inflation, fixes $A_s = H_I^4/(4\pi^2\dot{\phi}_0^2) \approx 2.1 \times 10^{-9}$ [4], which determines $\dot{\phi}_0 \approx (58.9H_I)^2$.

The above expression (5) characterizes the probability of particle production with physical momentum $k_p = k\eta_*H_I$. The total number density of particles can then be computed by integrating over all such k modes,

$$\begin{aligned}
 n &= \frac{1}{2\pi^2} \int_0^\infty dk_p k_p^2 e^{-\pi k_p^2 / (g|\dot{\phi}_0|)} e^{-\pi(M_0^2 - 2H_I^2) / (g|\dot{\phi}_0|)}, \\
 &= \frac{1}{8\pi^3} (g\dot{\phi}_0)^{3/2} e^{-\pi(M_0^2 - 2H_I^2) / (g|\dot{\phi}_0|)}. \quad (6)
 \end{aligned}$$

$$\begin{aligned}
 N_{\text{spots}} &= n \times \left(\frac{a_*}{a_0}\right)^3 \times 4\pi\chi_{\text{rec}}^2 \Delta\eta_s = \frac{1}{2\pi^2} \left(\frac{g\dot{\phi}_0}{H_I^2}\right)^{3/2} \frac{\Delta\eta_s}{\chi_{\text{rec}}} (k_*\chi_{\text{rec}})^3 e^{-\pi(M_0^2 - 2H_I^2) / (g|\dot{\phi}_0|)}, \\
 &\approx 4 \times 10^8 \times g^{3/2} \left(\frac{\Delta\eta_s}{100 \text{ Mpc}}\right) \left(\frac{100 \text{ Mpc}}{\eta_*}\right)^3 e^{-\pi(M_0^2 - 2H_I^2) / (g|\dot{\phi}_0|)}. \quad (7)
 \end{aligned}$$

Here a_* and $a_0 = 1$ are the scale factors at the time of particle production and today, respectively. The quantity χ_{rec} is the distance of the CMB surface from us and approximately equals 13871 Mpc, obtained from Planck's best-fit Λ CDM parameters, and $k_* = a_* H_I = 1/\eta_*$ is the mode that exits the horizon at the time of particle production.

B. Effect on the CMB

We now discuss the detailed properties of the spots and how they modify the CMB.

1. Primordial curvature perturbation from heavy particles

Owing to their large mass, the heavy particles can backreact on the spacetime metric around their locations, and can give rise to nontrivial curvature perturbations. The profile of such a curvature perturbation can be computed using the in-in formalism and the result is given by [26]

$$\langle \zeta_{\text{HS}}(r) \rangle = \frac{H_I}{8\epsilon\pi M_{\text{pl}}^2} \begin{cases} M(\eta = -r), & \text{if } r \leq \eta_* \\ 0, & \text{if } r > \eta_* \end{cases}. \quad (8)$$

Here $\epsilon = |\dot{H}_I|/H_I^2$ is a slow-roll parameter, and we have anticipated that this curvature perturbation would give rise to a hotspot (HS), rather than a coldspot. Importantly, the variation of the mass as a function of conformal time η controls the spatial profile. This variation can be computed from Eq. (4) by noting the slow-roll equation $\phi - \phi_* \approx \dot{\phi}_0(t - t_*)$, which gives

$$M(\eta)^2 = \frac{g^2 \dot{\phi}_0^2}{H_I^2} \ln(\eta/\eta_*)^2 + M_0^2. \quad (9)$$

Here we have used the relation between cosmic time t and the conformal time η , that also determines the size of the comoving horizon, $t - t_* = -(1/H_I) \ln(\eta/\eta_*)$.

Using the slow-roll relation $\dot{\phi}_0^2 = 2\epsilon H_I^2 M_{\text{pl}}^2$ and the fact that $M_0^2 \sim g|\dot{\phi}_0|$ so that N_{spots} is not significantly exponentially suppressed [see Eq. (7)], we can drop the

From an observational perspective, it is more convenient to relate n to the total number of spots that would be visible on the CMB sky. To that end, we need to specify the associated spacetime volume. Considering a shell of thickness $\Delta\eta_s$ around the CMB surface, the total number of spots in that shell is given by [27]

contribution of the second term in Eq. (9) away from η_* . The profile can then be simply written as

$$\langle \zeta_{\text{HS}}(r) \rangle = \frac{gH^2}{4\pi|\dot{\phi}_0|} \ln(\eta_*/r) \theta(\eta_* - r). \quad (10)$$

Given the typical size of a standard quantum mechanical fluctuation $\langle \zeta_q^2 \rangle^{1/2} \sim H^2 / (2\pi\dot{\phi}_0)$, we see the curvature perturbation associated with a hotspot differs primarily by $g/2$. In this work we will choose $g \sim \mathcal{O}(1)$, so the two types of perturbations will be of the same order of magnitude.

2. CMB anisotropy

After these fluctuation modes reenter the horizon, they source temperature anisotropies and give rise to localized spots on the CMB sky. To compute the resulting anisotropies, we first write metric perturbations,

$$ds^2 = -(1 + 2\Psi)dt^2 + a^2(t)(1 + 2\Phi)\delta_{ij}dx^i dx^j, \quad (11)$$

in the Newtonian gauge. The temperature fluctuations of the CMB corresponding to Fourier mode \vec{k} , pointing to direction \hat{n} in the sky is given by

$$\Theta(\vec{k}, \hat{n}, \eta_0) = \sum_l i^l (2l + 1) \mathcal{P}_l(\hat{k} \cdot \hat{n}) \Theta_l(k, \eta_0). \quad (12)$$

Here the multipole $\Theta_l(k, \eta_0)$ depends on the primordial perturbation $\zeta(\vec{k})$ and a transfer function $T_l(k)$ as

$$\Theta_l(k, \eta_0) = T_l(k) \zeta(\vec{k}), \quad (13)$$

with η_0 denoting the conformal age of the Universe today. Importantly, for our scenario $T_l(k)$ itself can be computed exactly as in the standard Λ CDM cosmology. It can be computed after taking into account the Sachs-Wolfe (SW), the integrated Sachs-Wolfe (ISW), and the Doppler (Dopp) effect [35],

$$\begin{aligned}\Theta_l(k, \eta_0) &\simeq [\Theta_0(k, \eta_{\text{rec}}) + \Psi(k, \eta_{\text{rec}})]j_l(k(\eta_0 - \eta_{\text{rec}})) + \int_0^{\eta_0} d\eta e^{-\tau} [\Psi'(k, \eta) - \Phi'(k, \eta)]j_l(k(\eta_0 - \eta)) \\ &\quad + 3\Theta_1(k, \eta_{\text{rec}}) \left[j_{l-1}(k(\eta_0 - \eta_{\text{rec}})) - (l+1) \frac{j_l(k(\eta_0 - \eta_{\text{rec}}))}{k(\eta_0 - \eta_{\text{rec}})} \right], \\ &\equiv [f_{\text{SW}}(k, l, \eta_0) + f_{\text{ISW}}(k, l, \eta_0) + f_{\text{Dopp}}(k, l, \eta_0)]\zeta(\vec{k}),\end{aligned}$$

where τ is the optical depth. The above expression relates a primordial perturbations ζ to a temperature anisotropy Θ_l .

3. Temperature anisotropy due to heavy particles

Regardless of the origin of $\zeta(\vec{k})$ is, we can compute $f_{\text{SW}}(k, l, \eta_0)$, $f_{\text{ISW}}(k, l, \eta_0)$, and $f_{\text{Dopp}}(k, l, \eta_0)$ as in the standard Λ CDM cosmology. Thus converting the position space profile in Eq. (10) to momentum space and using Eq. (II B 0 b), we can get the observed profile of a hotspot on the CMB sky. This Fourier transform of the profile (10) can be written as

$$\langle \zeta_{\text{HS}}(\vec{k}) \rangle = e^{-i\vec{k}\cdot\vec{x}_{\text{HS}}} \frac{f(k\eta_*)}{k^3}, \quad (14)$$

with a profile function

$$f(x) = \frac{gH^2}{\phi_0} (\text{Si}(x) - \sin(x)), \quad \text{Si}(x) = \int_0^x dt \sin(t)/t. \quad (15)$$

We parametrize the distance to the hotspot as

$$\vec{x}_0 - \vec{x}_{\text{HS}} = -(\eta_0 - \eta_{\text{HS}})\hat{n}_{\text{HS}}. \quad (16)$$

Here \vec{x}_0 and \vec{x}_{HS} parametrize our and the hotspot locations, respectively, and \hat{n}_{HS} points to the direction of the hotspot. The quantity η_{HS} denotes the location of the hotspot in conformal time with η_0 being the size of the present epoch. In the earlier paper, we took the hotspot to be on the CMB surface and hence set $\eta_{\text{HS}} = \eta_{\text{rec}} \approx 280$ Mpc. In this work, we allow the hotspots to be away from the last scattering surface with η_{HS} between $\eta_{\text{rec}} - \eta_*$ and $\eta_{\text{rec}} + \eta_*$, and study their signals on the CMB surface. This set up is summarized in Fig. 1.

As derived earlier, the temperature due to the hotspot is given by (dropping η_0 from the argument),

$$\Theta(\vec{x}_0, \hat{n}) = \int \frac{d^3\vec{k}}{(2\pi)^3} e^{i\vec{k}\cdot(\vec{x}_0 - \vec{x}_{\text{HS}})} \sum_l i^l (2l+1) \mathcal{P}_l(\hat{k} \cdot \hat{n}) (f_{\text{SW}}(k, l) + f_{\text{ISW}}(k, l) + f_{\text{Dopp}}(k, l)) \frac{f(k\eta_*)}{k^3}. \quad (17)$$

Here \hat{n} denotes the direction of observation. The functions $f_{\text{SW}}(k, l)$ and $f_{\text{ISW}}(k, l)$ are extracted from the transfer function using CLASS [36,37] as in Ref. [27]. Using the plane wave expansion

$$e^{-i\vec{k}\cdot\vec{r}} = \sum_{\ell=0}^{\infty} (-i)^\ell (2\ell+1) j_\ell(kr) \mathcal{P}_\ell(\hat{k} \cdot \hat{r}) \quad (18)$$

and the relation

$$\mathcal{P}_l(\hat{k} \cdot \hat{n}) = \frac{4\pi}{(2l+1)} \sum_{m=-l}^l Y_{lm}(\hat{n}) Y_{lm}^*(\hat{k}), \quad (19)$$

we get

$$\Theta(\vec{x}_0, \hat{n}, \eta_{\text{HS}}) = \frac{1}{2\pi^2} \int_0^\infty \frac{dk}{k} \sum_l j_l(k(\eta_0 - \eta_{\text{HS}})) (2l+1) \mathcal{P}_l(\hat{n} \cdot \hat{n}_{\text{HS}}) T_{\text{sum}}(k, l) f(k\eta_*), \quad (20)$$

$$T_{\text{sum}}(k, l) \equiv f_{\text{SW}}(k, l) + f_{\text{ISW}}(k, l) + f_{\text{Dopp}}(k, l). \quad (21)$$




FIG. 1. Representation of a hotspot on the CMB sky. Our location and the location of a hotspot are denoted as \vec{x}_0 and \vec{x}_{HS} , respectively, defined with respect to an arbitrary coordinate system. The black circle denotes the surface of last scattering, located at $\eta_{\text{rec}} \approx 280$ Mpc in conformal coordinates. Due to momentum conservation, heavy particles are produced in pairs, and the distance between the two members of a pair can vary between 0 and η_* . Therefore, in our analysis we allow the two members to be anywhere within the gray shaded region. We compute the temperature profile of a hotspot as a function of direction of observation \hat{n} , with the hotspot center in the direction of \hat{n}_{HS} .

Note $\Theta(\vec{x}_0, \hat{n}, \eta_{\text{HS}})$ depends on η_{HS} , the location of the hotspot—which need not be on the last scattering surface as mentioned above. Given the spherically symmetric profile of the hotspot, the Doppler contribution to $\Theta(\vec{x}_0, \hat{n}, \eta_{\text{HS}})$ is

$$\Theta_{\text{central}}(\vec{x}_0, \eta_{\text{HS}}) = \frac{1}{2\pi^2} \int_0^\infty \frac{dk}{k} \sum_l j_l(k(\eta_0 - \eta_{\text{HS}})) (2l+1) T_{\text{sum}}(k, l) f(k\eta_*). \quad (22)$$

In Fig. 2 we show the SW and ISW contributions to the central temperature as a function of η_{HS} after multiplying by the average CMB temperature $T_0 = 2.7$ K for $\eta_* = 160$ Mpc. For completeness, we also show the central temperature in Fig. 3, as obtained in [27], as a function of hotspot size η_* , assuming the hotspot is located on the surface of last scattering. As we can see, the pair-produced CMB spots are indeed hotspots when $\eta_* \lesssim \text{Gpc}$. For $\eta_* > 6600$ Mpc *coldspots* as opposed to hotspots arise. This is because the negative SW contribution dominates the positive ISW contribution, with the combination being negative.

III. SIMULATION OF THE CMB AND PHS SIGNALS

In order to design a PHS search, we simulate the PHS signal and CMB maps so that we can estimate the signal




FIG. 2. Central temperature $\Theta_{\text{central}} \times T_0$ of a hotspot as a function of the (radial) location of the hotspot. We choose $\eta_* = 160$ Mpc and $g = 1$. The dotted gray line indicates the location of the recombination surface. Larger (smaller) η_{HS} implies the hotspots are closer to (further from) us. We also show contribution of the Sachs-Wolfe term (orange) and the integrated Sachs-Wolfe term (purple) in determining the total temperature (olive). The left and right edges of the plot are at $\eta_{\text{HS}} = \eta_{\text{rec}} - \eta_*$ and $\eta_{\text{HS}} = \eta_{\text{rec}} + \eta_*$, respectively.

small, from now on we only include the SW and ISW corrections for our analysis.

4. Central temperature

It is useful to compute the temperature anisotropy at the central part of a hotspot. To that end, we set $\hat{n} = \hat{n}_{\text{HS}}$, implying $\mathcal{P}_l(\hat{n} \cdot \hat{n}_{\text{HS}}) = 1$, and

capture rate (“true positive rate”), and the background count for a CNN analysis. We notice that there are three types of backgrounds to consider for a PHS search: (i) the noise of the CMB detector, (ii) the astrophysical foreground, and (iii) the background from the standard primordial fluctuations.

A realistic analysis needs to take into account detector noise and foregrounds. In our analysis, we consider profiles on relatively large angular scales, $\ell < 1000$. For these scales current CMB temperature data, such as from Planck, is signal dominated and we thus do not need to add instrumental noise to our simulations. The astrophysical foreground comes from compact objects such as galaxies, galaxy clusters, gas, and dust, which can also produce localized signals. Part of these astrophysical foregrounds can be cleaned out due to their frequency dependence (for a




FIG. 3. Central temperature (green) of a hotspot originating from a heavy particle for $g = 1$, based on Eq. (22) with $\eta_{\text{HS}} = \eta_{\text{rec}}$. The green line illustrates the variation of the observed anisotropy as a function of the “size” of the hotspot, determined by the comoving horizon η_* at the time of particle production. The horizontal gray line gives a rough benchmark of the magnitude of the large-scale temperature anisotropy due to only the standard quantum fluctuations of the inflaton $(1/5)\langle \zeta_q^2 \rangle$, without taking into account acoustic oscillations. The dashed vertical gray lines show the benchmark choices for the hotspot size $\eta_* = 50, 100, 160$ Mpc chosen in the subsequent discussion. We take the plot from Ref. [27].

review see, e.g., Ref. [38]). For the signal sizes that we consider, corresponding to $\ell < 1000$, we do not expect significant astrophysical contamination after foreground cleaning and masking of the galactic plane, while for significantly smaller scales a detailed study of residual foregrounds and point sources would be required (see, e.g., Planck’s component separation analysis [39]). In the following, we therefore only consider the background from the primordial, almost Gaussian, fluctuations when studying the PHS signal. This last type of background is “irreducible” in the sense that it will always be present, originating from the fluctuations of the inflaton itself. We will assume the CMB maps are masked to reduce the astrophysical foregrounds and badly conditioned pixels and retain only 60% of the sky for the analysis. The number is similar to the sky fraction used in the Planck analysis [40].

Unlike the analysis in [27] that was based on a HEALPix [41] simulation, in this work, we use the QUICKLENS package² to simulate the CMB maps. QUICKLENS allows us to work in the “flat sky approximation,” neglecting sky curvature that is irrelevant to the size of the PHS profile we consider, as well as to draw sample maps with periodic boundary conditions to avoid complications due to masking. QUICKLENS can take a theoretical temperature power spectrum to produce mock flat sky CMB maps. To provide an initial input, we use the CLASS (v3.2) package [36,37] to compute a temperature anisotropy spectrum C_ℓ^{TT} based on the Planck 2018 [42] best fit Λ CDM parameters,

$$\{\omega_{\text{cdm}}, \omega_b, h, 10^9 A_s, n_s, \tau_{\text{reio}}\} \\ = \{0.120, 0.022, 0.678, 2.10, 0.966, 0.0543\}. \quad (23)$$

²<https://github.com/dhanson/quicklens>.

We will comment on the sensitivity of the CNN analysis to the Λ CDM parameters in Sec. IV A and Appendix A. We specify $\ell_{\text{max}} = 3500$ in the code for the maximum number of ℓ modes used for the image generation. As explained above, our signal profiles have support on length scales corresponding to an $\ell < 1000$, where instrumental noise is negligible compared to the primary background from CMB and can thus be ignored. An application to significantly smaller angular scales would need to take into account the noise properties of the experiment. We choose the image resolution such that 1 pixel = 10^{-3} radians to match Planck’s angular resolution down to ≈ 5 arc minutes [43]. We also use the relation between the angle and the comoving length on the last scattering surface $\Delta\eta/\chi_{\text{rec}}$.³ For instance, if the separation between two hotspot centers is 160 Mpc on the last scattering surface, then the two centers are 12 pixels away on the image, with $\chi_{\text{rec}} = 13871$ Mpc for Planck’s best-fit Λ CDM parameters.

For the CNN analysis, we begin by generating 360^2 pixel images, corresponding to a $[-10.32^\circ, 10.32^\circ]$ region in longitude and latitude ($n_x = 360$ in QUICKLENS). We then cut out a 90^2 patch from each of the 360^2 -sized maps. These nonperiodic, smaller maps are then used for further analysis. In particular, for our CNN analysis, we generate 160k training images, 40k validation 90^2 pixel images, and

³In Ref. [27], the angular size of one pixel was obtained by matching the pixel number to the total degrees of freedom in the ℓ modes ($\ell_{\text{max}}^2 + \ell_{\text{max}} = 4\pi/\theta_{\text{pixel}}^2$), together with the approximation $\ell_{\text{max}} \simeq \eta_0/\eta_{\text{pixel}}$. Although the matching reproduces the same angular resolution, the relation between ℓ_{max} and η_{pixel} gives $\Delta\theta = \sqrt{4\pi}\Delta\eta/\chi_{\text{rec}}$. Since $\ell_{\text{max}} \simeq \eta_0/\eta_{\text{pixel}}$ comes from the approximation of the k -mode integral with $j_\ell(k\chi_{\text{rec}})$ and $k = 2\pi/\eta$, the relation between the angle and length is less robust than $\Delta\theta = \Delta\eta/\chi_{\text{rec}}$.

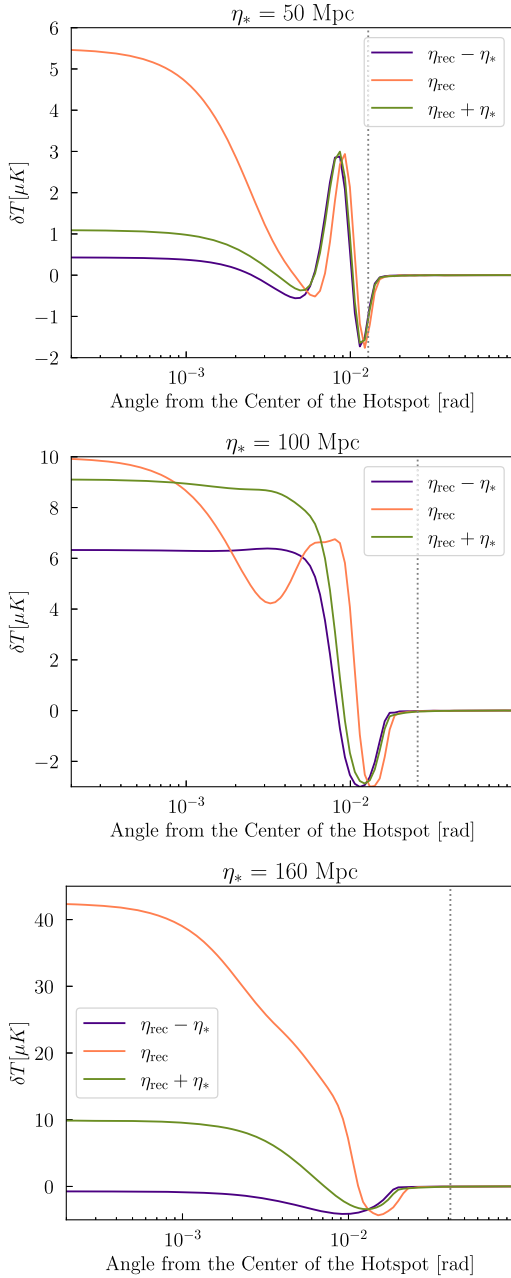


FIG. 4. Radial profile of a single hotspot with the heavy particle position inside (olive), on (orange), and outside (purple) of the last scattering surface. The locations of these hotspots in conformal time are taken to be $\eta_{\text{rec}} + \eta_*$, η_{rec} , and $\eta_{\text{rec}} - \eta_*$, respectively, as denoted by the labels. From upper left to bottom: horizon size for the hotspot production at $\eta_* = 50, 100, 160$ Mpc. The plots assume the inflaton- χ coupling $g = 1$.

an additional 5k test images to quantify the network performance. Training the neural network on smaller patches yields better training convergence and does not lead to loss of information as long as the characteristic size of the signal is smaller than the size of the patch.

The profile of each of the PHSs is described by Eq. (20), where the function depends on the distance to the hotspots

($\eta_0 - \eta_{\text{HS}}$) and the angle $\cos^{-1}(\hat{n} \cdot \hat{n}_{\text{HS}})$, as defined in Fig. 1. The overall magnitude of the signal temperature is proportional to the coupling g . When generating the signal, we require both the hotspots to be within a shell $\pm\eta_*$ around the last scattering surface as shown in Fig. 1. For example, when studying the case with $\eta_* = 160$ Mpc, we first divide the ± 160 Mpc region into 50 concentric annuli, each having equal thickness. We then choose the first hotspot from a pair to lie on any of these 50 annuli with equal probability. The second member is then chosen anywhere within a sphere of radius η_* centered on the first hotspot, again with a uniform random distribution.⁴ A pair is kept for further analysis only if both the spots of the pair falls within the $\pm\eta_*$ shell of the last scattering surface. Since the distribution in a 3D volume allows hotspots to orient along the line-of-sight direction, the average separation between the two hotspots projected on the last scattering surface is smaller than the separation assumed in Ref. [27] that only considered PHSs on the last scattering surface.

Once we generate PHS images with random orientation and separation between two hotspots, we pixelate them and add the PHS image to the simulated CMB maps to produce the signal image. We follow this procedure for all the signal images in our study. In this work, we study benchmark models with horizon sizes

$$\eta_* = 50, 100, 160 \text{ Mpc}, \quad (24)$$

and couplings from $g = 1$ to 4. Specifying g and η_* sets the overall temperature and the profile of the hotspot, *à la* Eq. (20). Within the approximations we have made in Sec. II, the remaining model parameter, M_0 , only affects the overall number of hotspots N_{PHS} [through Eq. (7)]. Going forward, we will compute the number of hotspots that can be hidden within the background fluctuations for given benchmark coupling and η_* . Then, using Eq. (7), the upper bounds on N_{PHS} can be translated into lower bounds on M_0 . As an illustration of what a benchmark PHS looks like, in Fig. 5 we show examples of the CMB background (left), PHS signal (middle), and the signal plus background (right) for $g = 4$ with different choices of η_* . Note that it is difficult to identify the signals by eye in the plots on the right, even with such a large coupling.

Compared to Ref. [27], the benchmark η_* values are identical, but we choose smaller values of the coupling g . This is because we find the CNN analysis is much more powerful than the “cut and count” method adopted in

⁴Our motivation for the uniform distribution is driven by Eq. (6). There, the integral is dominated by $k_p \sim M_0$, as small k_p are suppressed by the k_p^2 factor and large k_p are exponentially squashed. Assuming that χ are produced semirelativistically and can travel a horizon scale distance, we expect a uniform distribution of hotspots within $\pm\eta_*$ of the last scattering surface.




FIG. 5. Example plots of pure background from QUICKLENS simulation (left), pure signals (middle), and signals with $g = 4$ on top of the simulated background (right). The scalar particles are produced at comoving horizon sizes $\eta_* = 50$ Mpc (top), 100 Mpc (middle), $\eta_* = 160$ Mpc (bottom). The signals at different benchmark η_* have roughly the same size, as the η_* dependence only enters logarithmically. The two hot spots are clearly separated for $\eta_* = 160$ Mpc and $\eta_* = 100$ Mpc, while for $\eta_* = 50$ Mpc they overlap.

Ref. [27], and therefore capable of identifying fainter hotspots. We chose the benchmark η_* values to test out a variety of different PHSs; $\eta_* = 160$ Mpc hotspots have a very high central temperature (Fig. 3), while $\eta_* = 50$ Mpc hotspots are significantly cooler and have smaller interspot separation. The choice $\eta_* = 100$ Mpc sits between these for comparison.

IV. IDENTIFYING PAIRWISE HOTSPOTS WITH CNN

In this section we describe the training process for the CNN using 90^2 pixel images, and discuss some qualitative properties of the training result. We then apply the trained network to a larger sky map and present results on the upper bound on the number of PHSs for given values of η_* and g . We end the section with some comparisons between the CNN and a matched filter analysis.

A. Network training on small sky patches

CNNs are one of the most commonly used deep neural networks specialized for image recognition [44,45]. In this study, we build the network using PyTorch [46] with the structure shown in Fig. 6. The network takes a CMB or CMB + PHS image as an input and outputs a single value between 0 and 1, which can be interpreted as the probability of the input image containing the PHSs. We train the network on 160k images (see Sec. III), half of which contain a single pairwise hotspot profile on top of the CMB and the rest are CMB-only images. For optimization, we use a binary cross entropy loss function, commonly used for binary classification, along with Adam optimizer [47] and 10^{-4} learning rate. We briefly experimented with variations of the network architecture, in particular by changing the number of convolutional filters and kernel sizes. We did not observe a significant change in network performance, most likely indicating that our default network is sufficiently expressive to solve the task optimally.

We train the network using PHS signals with $g = 3$ for all the three values of η_* individually. One may wonder how well a network trained on one g value will generalize to different values without retraining. As the CNN (unlike the matched filter discussed below) is nonlinear, extrapolation to values of g other than what was used for training is not guaranteed to be optimal. On the other hand, training a




FIG. 6. A schematic architecture of the CNN used in this work. We applied two convolutional layers in series; first, eight kernels with size of 16×16 and stride of 2 are applied, then, eight independent kernels size of 8×8 yields feature map of 40. Next, we apply a max pooling using the kernel and stride size of 2×2 , which subsequently reduces the image dimension down to $20 \times 20 \times 8$. Processed images get further reduced by going through 2D convolution and max pooling, further reducing the size of the image to $5 \times 5 \times 8$. After four sets of total convolution followed by average pooling, the final feature maps are flattened to feed into fully connected network form, and the final network ends with single output value that sits between 0 and 1. Throughout the network, we use the rectified linear unit function [48] to introduce nonlinearity, except for the output layer, which has a sigmoid activation function suitable for the binary classification.

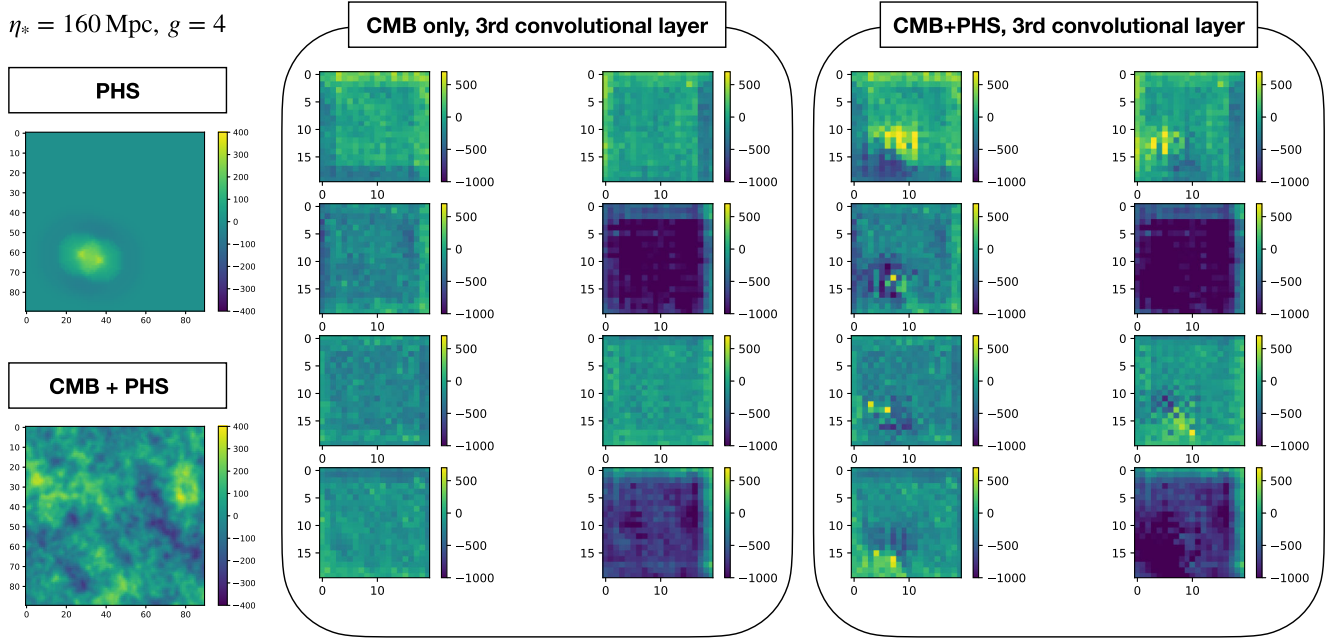


FIG. 7. Comparison between true and the CNN feature maps with and without implanted signals. The left plots show the PHS signal and signal plus the CMB background. The middle and right plots show feature maps after going through three convolutional layers. The enhanced signal locations on the feature maps on the right align with the true location of the hotspots after rescaling the pixel coordinates with respect to the relative size between the third layer (20^2 pixels) and the original image (90^2 pixels). Here we take $\eta_* = 160$ Mpc, $g = 4$, and $\eta_{\text{HS}} = \eta_{\text{rec}}$ for both the spots.

CNN for each possible benchmark input is time and resource intense. Empirically, we find that the network trained at $g = 3$ works well over a wide range of g values, perhaps because the network learns to analyze the shape rather than the amplitude of the profile. In a fully optimal analysis one would want to retrain the neural network over a grid of g values.

To get some idea for how the CNN discriminates between signal and background images, we show the feature maps from the first three convolutions in Fig. 7 for $\eta_* = 160$ Mpc and $g = 4$. As we can see proceeding from left to right, the trained network does amplify the signal region compared to the background-only image, and the convolutional layers can emphasize the correct locations of each spot in the feature map.

To quantify the performance of the CNN, we generate a test sample of 5k CMB-only maps and 5k CMB + PHS maps, each having 90^2 pixels. For a CMB + PHS map, we inject one randomly oriented and located PHSs in the CMB map. The PHS signal occupies $\mathcal{O}(50^2)$ pixels in the examples that we study, and thus the 90^2 pixels image is only slightly larger than the signal. When an image has network output > 0.5 , we count it as an identified signal map. We call the signal capture rate (true positive rate, $\epsilon_{S,90^2}$) as the fraction of CMB + PHS images being correctly identified as signal maps, and define the fake rate (false positive rate, $\epsilon_{B,90^2}$) as the fraction of CMB-only images being wrongly identified as signal maps,⁵

$$\begin{aligned} \epsilon_{S,90^2} &= \frac{\text{number of signal-injected images with CNN output} > 0.5}{\text{total number of signal-injected images}}, \\ \epsilon_{B,90^2} &= \frac{\text{number of background-only images with CNN output} > 0.5}{\text{total number of background-only images}}. \end{aligned} \quad (25)$$

⁵In the actual search, there can be more than one PHS in a 90^2 -pixel region, and the CNN would still count the region to be one signal map. We verify that the signal capture rate would increase if there are more PHSs in the image. When we study the sensitivity of the CNN search, having additional PHSs around the same location will help the search, and this makes our analysis based on having one PHS in a 90^2 -pixel image to be conservative. Moreover, given that the CNN search can probe PHSs with a small number of signals on the CMB sky, the probability of having additional PHSs around the same location is small. Therefore, counting the number of 90^2 -pixel regions should give a good approximation of the PHSs in the analysis.

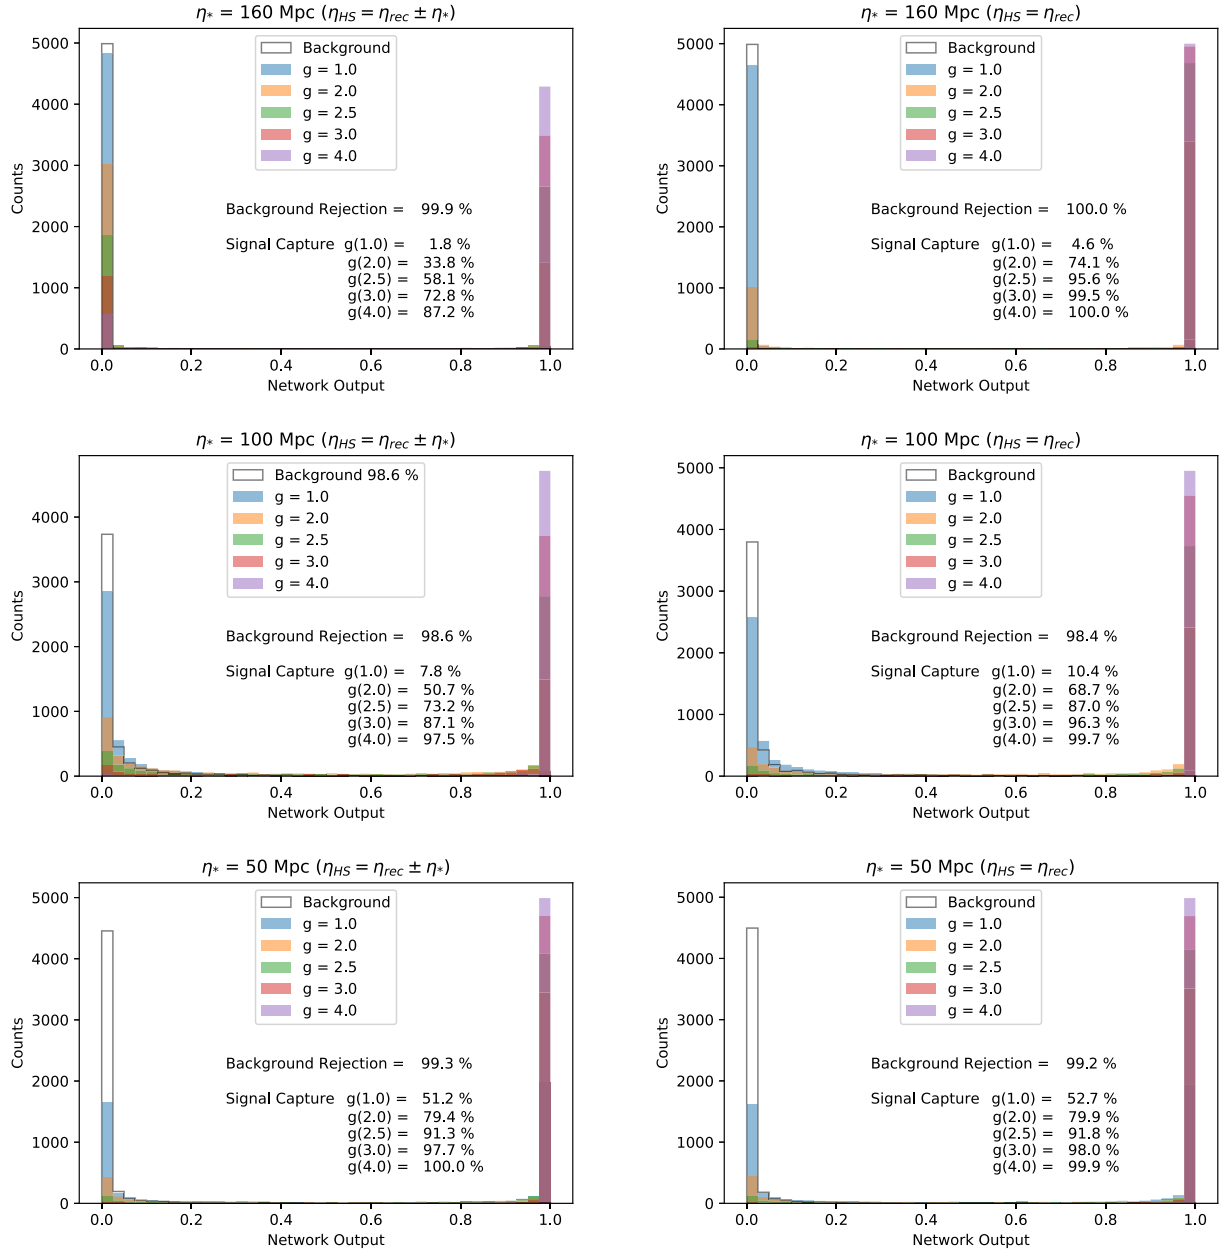


FIG. 8. Network output for 5k images without (blank histogram) and with (colored histograms) PHS signals. We count the image as an identified signal map when the network output > 0.5 . In the plots we show the background rejection rate from the CMB-only analysis and the signal capture rate from the CMB + PHS images, for different inflaton- χ couplings g . The fake rate is defined as $(1 - \text{background rejection rate})$. The plots on the left have both the hotspots distributed uniformly with separation $\leq \eta_*$ and within $\eta_{HS} = \eta_{rec} \pm \eta_*$, which is how we simulate the signal for the rest of the study. The signal capture rate therefore includes possible suppression due to hotspots moving off the last scattering surface. For comparison, we show the training results in the right plots requiring $\eta_{HS} = \eta_{rec}$. Comparing results obtained from the same study but with different sets of 5k images, we find the efficiency numbers vary by $\sim 0.1\text{--}1\%$.

In Fig. 8, we show the network output for the 5k images with and without injecting the PHS signal. In the left column we show the result when the PHSs are uniformly distributed within a shell of $\eta_{rec} \pm \eta_*$ around the surface of last scattering, while the right column shows the result when $\eta_{HS} = \eta_{rec}$. The signal capture and background rejection rates in Fig. 8 refer to $\epsilon_{S,90^2}$ and $(1 - \epsilon_{B,90^2})$.

Clearly, for $g \geq 3$, our CNN setup is highly efficient at separating CMB + PHS images from CMB images alone. For example, for $g = 3$ (the same coupling as in the training sample) and $\eta_* = 160$ Mpc, $\epsilon_{S,90^2}$ is over 73% with $\epsilon_{B,90^2}$ less than 0.1%. For $\eta_* = 160$ Mpc and 100 Mpc, the signal capture rate falls if the hotspots are off the last scattering surface but in the $\eta_{rec} \pm \eta_*$ window we consider. When




FIG. 9. Left: PHS signals that are implanted on CMB map. Right: probability map from scanning the same 720^2 image plus the CMB with the CNN search of 90^2 -pixel region shifting in steps of 1 pixel. The true and fake signals show up as clusters in the processed image. We further suppress the number of fake signals in the following analysis by applying cuts on the network output of each pixel and the pixel number in each cluster. We find the analysis from shifting the search window in steps of 5 pixels produce similar results to the steps of 1 pixel and therefore use 5-pixel steps for the rest of the analysis.

applying the same trained network on dimmer PHS signals ($g < 3$), $\epsilon_{S,90^2}$ drops, but the background rejection rate remains close to unity.

Both $\epsilon_{S,90^2}$ and $\epsilon_{B,90^2}$ vary with the horizon size. Comparing results for $\eta_* = 160$ Mpc to $\eta_* = 50$ Mpc, the $\epsilon_{S,90^2}$ values are similar for $g \geq 3$, but $\eta_* = 50$ Mpc case performs much better at weaker coupling ($\epsilon_{S,90^2} = 51.2\%$ for $\eta_* = 50$ Mpc compared to 1.8% for $\eta_* = 160$ Mpc, both for $g = 1$). The $\eta_* = 50$ Mpc case has a larger background fake rate, compared to $\eta_* = 160$ Mpc. However, even if we incorporate the background and compare $\epsilon_{S,90^2} / \sqrt{\epsilon_{B,90^2}}$ —the efficiency ratio is $\mathcal{O}(10)$ times larger for the dimmer, $\eta_* = 50$ Mpc case. The ability of catching dimmer signals indicates that the network uses additional information than the overall temperature to identify the PHSs.

Although it is difficult to know exactly how the CNN identifies the PHSs, the network seems to more accurately identify PHSs with a distinct rim structure compared to just utilizing the fact that there are two hotspots (Fig. 5). One indication that the CNN utilizes the rim structure of the $\eta_* = 50$ Mpc signal is that the signal capture rate for that benchmark is insensitive to whether or not the PHSs lie on the last scattering surface. We perform the same CNN analysis by having the signal hotspots centered on the last scattering surface [$\eta_{\text{HS}} = \eta_{\text{rec}}$ in Eq. (20)] and summarize results in the right column of Fig. 8. For hotspots with temperature profile peaked at center, as we show in the $\eta_* = 160$ and 100 Mpc plots in Fig. 4, the highest PHS temperature takes the maximum value when $\eta_{\text{HS}} = \eta_{\text{rec}}$ (orange). It then is reasonable to have a larger average signal capture rate when the hotspots center on the last scattering surface. However, as we illustrate in the upper left plot in Fig. 4, the “shell” of the $\eta_* = 50$ Mpc signal in 3D always project into a rim with a fixed temperature (at angle ≈ 0.008 rad), regardless of the location of the hotspot, $\eta_{\text{rec}}, \eta_{\text{rec}} + \eta_*$, or $\eta_{\text{rec}} - \eta_*$. Therefore, if the CNN identifies

the $\eta_* = 50$ Mpc signal based on the rim structure, $\epsilon_{S,90^2}$ should remain the same even when the PHSs are on the last scattering surface. This is indeed what we see on the bottom plots in Fig. 8. Further study on what features the CNN uses to identify the $\eta_* = 50$ Mpc case can be found in Appendix C.

B. Application of the trained network to larger sky maps

After training the CNN to identify PHSs in images with 90^2 pixels, we look for signals on a larger sky map by applying the same network analysis repeatedly across the larger map. In this way we can analyze, in principle, arbitrarily large maps. A benefit of such a larger map search is that it avoids the loss of sensitivity to signals where a PHS is partially cut out by the boundary of a 90^2 -pixel region. Such a PHS would be lost had we simply partitioned the sky into nonoverlapping 90^2 -pixel regions.

For a concrete application, we study maps with 720^2 pixels⁶ using the following steps: (i) we apply the trained network on the upper left corner of the map, obtaining the network output, (ii) we shift the 90^2 -pixel “window” to the right by five pixels and get the network output again, (iii) repeat the process until we hit the right-hand side of the large map. Then, return to the upper left corner but slide the window down by five pixels, (iv) continue with these steps until the entire larger map is covered. The result of steps (i)–(iv) result in what we call a “probability map.” Starting with an original 720^2 image and scanning in steps of five pixels, the probability map has 126^2 entries, with each entry

⁶Repeating this analysis on even larger maps would be ideal and lessen the assumptions made when extrapolating to the whole sky. However, practically, we found that 720^2 was the largest size we could make without sacrificing statistics (number of maps).

showing the probability of having a signal in a 90^2 -pixel region centered at each pixel. We have tried different step sizes and find that a 5-pixel step size yields nearly identical results to a 1-pixel step size for the following analysis, so we use the 5 pixel step size for improved computational speed.

In Fig. 9, we show an example of the probability map (right) obtained from the image of CMB plus three PHS signals shown in the left plot. To make the signals more visible by eye, we make the probability map by sliding the search window in 1 pixel step, thereby giving a 630^2 -pixel map. The true signals in the right plot show up as three bright clusters, and there are fake signals from the CMB fluctuations themselves. To reduce these fake signals, we further apply cuts on the probability map by only keeping pixels satisfying a threshold cut (network output) > 0.9 .

To properly count the number of observed signals, we cluster nearby pixels in the probability map. Specifically, we employ the ‘‘scikit-image’s morphology label’’ technique [49] to connect neighboring pixels with the same values and therefore group the connected pixels into clusters. We further remove clusters with < 30 connected pixels, as these have a smaller chance to be a true signal. The choices of cuts on network output threshold and number of pixels in a cluster were made by trying several values and choosing the value that maximized the signal capture rate while keeping the fake rate small.

To determine how well the method works, we study a set of 500 CMB-only images (i.e., pure background) and 500 images where $N_{\text{inj}} = 3$ randomly distributed PHS signals have been injected. We then define ‘‘efficiencies’’:

$$\begin{aligned} \epsilon_{S,720^2} &= \frac{\text{total number of clusters passing the threshold in CMB + PHS maps}}{3 \times 500}, \\ \epsilon_{B,720^2} &= \frac{\text{total number of clusters passing the threshold in CMB-only maps}}{500}. \end{aligned} \quad (26)$$

The results for the benchmark η_* values and three different couplings are summarized in Table I below. The 720^2 search retains some of the key features from the 90^2 search, namely the superior performance for $\eta_* = 50$ Mpc when g is small.

C. Obtaining theoretical bounds from detection statistics

Using Table I we can calculate the upper bound on N_{PHS} , the number of hotspot pairs produced where both members of a pair lie within a shell of $\pm\eta_*$ from the last scattering surface. As an example, let us take $\eta_* = 50$ Mpc and $g = 1$. From Table I, we see $\epsilon_{S,720^2} = 54.6\%$ while $\epsilon_{B,720^2} = 1.4\%$. Assuming that only a fraction $f_{\text{sky}} = 60\%$ is used for the search, the total number of signals for this benchmark is

TABLE I. CNN result from scanning 500 randomly generated CMB or CMB + PHS maps using the network trained in Sec. IV A. The image size is 720^2 pixels, and we shift the search window having 90^2 pixels by 5-pixel steps. The fake rate is the average number of fake signals from a 720^2 -pixel map with CMB only. The signal capture rate is the chance of identifying each input PHS signal. Comparing results obtained from the same study but with different sets of 500 images, we find the efficiency numbers vary by ~ 0.1 – 1% .

	$\eta = 50$ Mpc	$\eta = 100$ Mpc	$\eta = 160$ Mpc
$\epsilon_{B,720^2}$	1.4%	11%	6.6%
$\epsilon_{S,720^2}, g = 1$	54.6%	0.8%	0.5%
$\epsilon_{S,720^2}, g = 2$	84.0%	34%	34.6%
$\epsilon_{S,720^2}, g = 3$	98.6%	76.8%	71.2%

$\text{Sig} = \epsilon_{S,720^2} N_{\text{PHS}} f_{\text{sky}}$, while the number of background events is $Bg = 25\epsilon_{B,720^2} f_{\text{sky}}$, where the factor of 25 is the number of 720^2 patches needed to cover the full sky. From the number of signal and background events, we form the log-likelihood ratio [50,51] and then solve for N_{PHS} for the desired signal significance. When calculating the 2σ exclusion bound, we require

TABLE II. Upper: 2σ upper bound on the number of PHSs in the whole CMB sky with both hotspot centers located within $\eta_{\text{rec}} \pm \eta_*$ window around the last scattering surface. In the calculation we assume sky fraction $f_{\text{sky}} = 60\%$. Lower left: lower bounds on the bare mass of the heavy scalar field in units of the Hubble scale during the inflation. Lower right: lower bounds on the bare mass in units of the rate of the mass, $(g\dot{\phi}_0)^{1/2}$, owing to the inflaton coupling.

Number of PHSs	$\eta = 50$	$\eta = 100$	$\eta = 160$
$g = 1$	8	840	1162
$g = 2$	5	20	17
$g = 3$	4	9	8
M_0/H_I	$\eta = 50$	$\eta = 100$	$\eta = 160$
$g = 1$	145	120	114
$g = 2$	213	199	194
$g = 3$	266	253	247
$M_0/(g\dot{\phi}_0)^{1/2}$	$\eta = 50$	$\eta = 100$	$\eta = 160$
$g = 1$	2.5	2.0	2.0
$g = 2$	2.6	2.4	2.4
$g = 3$	2.6	2.5	2.4

TABLE III. Same as Table II but for the 5σ discovery reach.

Number of PHSs	$\eta = 50$	$\eta = 100$	$\eta = 160$
$g = 1$	16	2047	2757
$g = 2$	10	48	40
$g = 3$	9	21	19
M_0/H_I	$\eta = 50$	$\eta = 100$	$\eta = 160$
$g = 1$	143	116	110
$g = 2$	210	194	189
$g = 3$	262	247	241
$M_0/(g\dot{\phi})^{1/2}$	$\eta = 50$	$\eta = 100$	$\eta = 160$
$g = 1$	2.4	2.0	1.9
$g = 2$	2.5	2.3	2.3
$g = 3$	2.6	2.4	2.4

$$\sigma_{\text{exc}} \equiv \sqrt{-2 \ln \left(\frac{L(\text{Sig} + Bg|Bg)}{L(Bg|Bg)} \right)} \geq 2,$$

$$\text{with } L(x|n) = \frac{x^n}{n!} e^{-x}. \quad (27)$$

Note that this is the expected bound, as we are taking simulated CMB background to be the number of observed events [n in Eq. (27)]. The resulting values of N_{PHS} are given in the left panel of Table II. It is also interesting to determine how many PHSs would be needed for discovery at each benchmark point. We calculate the expected discovery reach using

$$\sigma_{\text{dis}} \equiv \sqrt{-2 \ln \left(\frac{L(Bg|\text{Sig} + Bg)}{L(\text{Sig} + Bg|\text{Sig} + Bg)} \right)} \geq 5. \quad (28)$$

The results are collected in Table III.

We can further obtain the minimum mass M_0 of the heavy particle corresponding to σ_{exc} and σ_{dis} using Eq. (7) and $\Delta\eta = 2\eta_*$.⁷ In Table II and III, we show the bounds (or reach) on the number of PHSs and M_0/H_I . Due to the

⁷One subtlety in solving the mass bound is that when simulating the PHS signals, we require both hot spots to be within $\pm\eta_*$ around the last scattering surface. Hence, the simulation excludes PHSs with one of the hot spots outside of the shell region that would be harder to see by the CNN. However, when solving the upper bound on the PHS density using Eq. (7), we take into account the signals that are partially outside of the shell region, leading to an overestimate of the signal efficiency and a stronger upper bound on the number density. From checking the hot spot distribution numerically, we find that $\approx 17\%$ of the PHSs in our examples can be partially outside of the $\pm\eta_*$ region. Fortunately, since the size of M_0 only depends on the number density bound logarithmically, the error only changes the M_0 bound by up to 1%. This is acceptable for the accuracy we want for the concept study.

energy injection from the dynamics of the inflaton, we can probe scalar particles with masses up to $\approx 260H_I$. In the bottom right tables, we show that the mass bounds correspond up to ≈ 2.6 times the mass-changing rate caused by the inflaton rolling ($\sqrt{g\dot{\phi}_0}$), which dominates the exponential suppression in Eq. (7). We also plot the 2σ lower bound on M_0/H_I in Fig. 10. Since the N_{PHS} depends on M_0 exponentially, a slightly lower scalar mass than the 2σ bound leads to a 5σ discovery of the PHSs.

These bounds are significantly improved compared to the previous analysis in Ref. [27]; this is not surprising given that the analysis in Ref. [27] was very simplistic, utilizing only a single temperature cut to separate signal from background. Using the CNN, we can now obtain meaningful bounds for $g = 1, 2$ —cases for which the PHSs were rather invisible before. For hotter signals, e.g., $g = 3$, the CNN analysis beats the past result by $\Delta M_0 \approx 60H_I$. This is a notable improvement given that the PHS density is exponentially sensitive to the scalar mass (squared).

Finally, to show that the CNN search of localized objects gives a better probe of heavy particle production than the measurement of CMB temperature power spectra, we plot the corrections to the $\Lambda\text{CDM } D_\ell^{\text{TT}}$ spectrum in Appendix B, including the same number of PHSs in Table II. For example, for $g = 1$, $\eta_* = 160$ Mpc, we see from Table II that the 2σ bound on N_{PHS} from our CNN analysis is 1162 hotspot pairs. Injecting 1162 hotspots into the sky,⁸ we find a correction to D_ℓ^{TT} of $\Delta\chi^2 = 0.3$ —which is well within the 1σ band on Planck 2018 temperature power spectrum. Repeating this exercise with the other benchmarks in Table II, yields $\Delta\chi^2$ values that are even smaller.

D. Comparison with a matched filter analysis

Matched filter analysis is a standard tool for identifying localized signals on a CMB map. Given a 2D power spectrum of the CMB, $P(k)$, we can obtain a filtered map $\psi(\vec{r})$ in position space from a convolution between the original image (signal plus background) $\zeta(\vec{k})$ and a signal filter $h(\vec{k})$ [the Fourier transform of a profile $h(\vec{r})$ in position space],

$$\psi(\vec{r}) = \int \frac{d^2\vec{k}}{(2\pi)^2} \left(\frac{\zeta(\vec{k})h(\vec{k})}{P(k)} \right) e^{i\vec{k}\cdot\vec{r}}. \quad (29)$$

If the signal is spherically symmetric, then the filter simplifies to $h(\vec{k}) = h(k)$. From the filtered map $\psi(\vec{r})$ one can construct an optimal likelihood ratio test between the Gaussian null hypothesis and the existence of the signal

⁸For simplicity, we restrict all hotspots to the last scattering surface. This somewhat overemphasizes the PHS correction to the power spectrum, as scenarios with both particles fixed to the last scattering surface are, on average, brighter than when η_{HS} varies.




FIG. 10. Bound on the heavy scalar mass for $\eta_* = 50$ Mpc, $\eta_* = 100$ Mpc, and $\eta_* = 160$ Mpc. In the region above the “% backreaction” line, the backreaction to the inflationary dynamics due to particle production is smaller than a percent (see Ref. [27] for a more detailed discussion). The light blue lines show various contours of N_{PHS} . We notice that the projected CNN search is able to cover most of the parameter space up to the target $N_{\text{PHS}} = 1$ contour.

(see, e.g., [24]), making the matched filter ideal for picking out single (or more generally, nonoverlapping) localized signals.

As we have seen, while the individual hotspots are spherically symmetric, they often overlap (at least for the range of parameters we are interested in), leading to a net signal in the sky that is no longer spherical. Additionally, the random separation between the initial heavy particles means the resulting PHSs are not uniform. The unusual shape and variability among signals make the PHSs less suitable for a vanilla matched filter analysis. While it may be possible to design a complicated and large bank of matched filters to cover the space of possible signal templates, the CNN analysis can effectively learn a set of flexible filters to enhance the signal over background even with varying and nonspherical signal shapes.

Even if the matched filter analysis defined in Eq. (29) is not optimal for the full pairwise hotspot signal, it is still instructive to compare a few examples of the matched filter analysis versus the CNN. For this comparison, we consider PHSs that lie only on the last scattering surface. The combined signal from the PHSs will still be nonspherical, but restricting all PHSs to the last scattering surface does take away some of the variability among signals.⁹ While each hotspot in a pair will “pollute” the other—meaning that it appears as a background that is different from the CMB fluctuations—each of the two hotspots can still be picked up effectively by the single spot template $h(k)$.

We perform the comparison using 90^2 pixel images with one PHS injection. We use QUICKLENS to generate the CMB maps, which follows periodic boundary condition and thereby ensures the separation between k modes in the 2D power spectrum $P(k)$ of the CMB image. The CNN

⁹We still allow a random separations (within η_*) between hotspots on the (2D) last scattering surface.

results for this signal set have already been shown in Sec. IV A and can be found in the right hand panels of Fig. 8; the background rejection is above 99% for all benchmark points, while the signal capture rate varies from a few percent to 100% depending on η_* and g .

For the matched filter analysis, we obtain $P(k)$ from the average of the discrete Fourier transform of 500 simulated images. We also apply discrete Fourier transform on the profile of a single hotspot in the PHSs, and use it as $h(k)$ in the convolution. Carrying out the integral in Eq. (29), we obtain the processed maps $\psi(\vec{r})$. An example of the signal processing is shown in Fig. 11, where the plot on the left is the PHS signal ($\eta_* = 160$ Mpc and $g = 2$), the middle is the signal plus background, and the right plot is the output image $\psi(\vec{r})$. We see that the filter can indeed pick up the signal hidden inside the background.

As one way to quantify the matched filter results, in Fig. 12 we show the distribution of largest $\psi(\vec{r})$ values in each of the 500 maps generated with (blue) and without (red) PHS signals with $\{\eta_*, g\} = \{160 \text{ Mpc}, 2\}$ (left) and $\{100 \text{ Mpc}, 2\}$ (right). From this perspective, the matched filter clearly separates the signal and background for the two cases. We also perform the same analysis for the $\eta_* = 50$ Mpc signals (which have much lower temperatures). In this case, the overlap between signal and background in the ψ distribution is large, and a simple ψ cut is not the optimal way to separate the signal and background. For this reason, we only consider the $\eta_* = 100$ and 160 Mpc examples in the following discussion.

To provide a rough numerical comparison between the matched filter and the CNN analysis, we apply a ψ_{max} cut in each of the matched filter histograms in Fig. 12. We choose the ψ_{max} cut value to equal the background rejection rate in the CNN analysis, then compare signal capture rates in the two analyses. For the $\eta_* = 160$ Mpc example, the signal capture rate is about 5% and 74% for $g = 1$ and 2, while the




FIG. 11. Example images from the matched filter analysis. Left: PHSs with $\eta_* = 160$ Mpc and $g = 2$. Middle: signal plus the background. Right: filtered map from the convolution integral Eq. (29).




FIG. 12. Maximum pixel distribution in filtered maps, where the value ψ of the pixels on the filtered map is defined in Eq. (29). We use 500 CMB-only and 500 CMB + PHS maps and plot the distribution of the maximum ψ of each filtered map to show the separation between the CMB and CMB + PHS results. We used feature scaling also known as min-max normalization for ψ_{\max} , so that the smallest value is zero and the largest value is 1.

match filter analysis performs slightly better, capture rates 8% and 98%, respectively. For $\eta_* = 100$ Mpc, the CNN signal capture rates are $\sim 10\%$ and $\sim 69\%$ for $g = 1$ and 2, while the match filter analysis rates are slightly lower, 4% and 50%. We note that the neural network was trained for $g = 3$ so the signal capture rate may be somewhat sub-optimal for the $g = 1$ and $g = 2$ cases. This could be alleviated by retraining.

In summary, the matched filter and neural network analysis perform similarly. This indicates that the neural network training has converged to a good solution that is not easily outperformed by the standard matched filter method. Our comparison is not exactly one-to-one because the neural network is trained on pairs of hotspots while the matched filter uses a single spherically symmetric profile. This gives the neural network an advantage in the realistic situation of pairs of hotspots. A full comparison of the two

methods would require an extended and complicated matched filter analysis which includes pairs of profiles.

V. DISCUSSION AND CONCLUSION

In this work, we show that CNNs provide a powerful tool to identify PHSs on the CMB sky. These PHSs can originate from superheavy particle production during inflation. We improve the previous analysis of Ref. [27] by more accurately modeling the distribution of PHSs on the CMB sky and by developing a CNN-based signal search strategy.

To accurately model the PHS distribution, we include the possibility that PHSs are distributed along the line-of-sight direction, rather than fixed to the last scattering surface. As a result, the average interspot separation within a PHS, when projected onto the CMB, is smaller than in Ref. [27].

For PHSs with small values of η_* , such as $\eta_* = 50$ Mpc, the two hotspots in a PHS significantly overlap with each other, and the resulting PHSs look like a single object, but with a distinct angular profile (Fig. 5).

For the signal search, we construct a CNN to identify PHSs from within the CMB, the standard fluctuations of which act as backgrounds for the signal. The network is trained on 90^2 pixel images with and without PHSs injected in them (both with hotspots distributed in 3D, and with hotspots fixed on the last scattering surface). During training we choose a coupling $g = 3$, but the trained CNN can still identify PHSs for smaller values of g with a significant signal capture rate and small background fake rate. We find that the CNN actually performs better for the smaller η_* benchmark, even though the hotspots are dimmer. We believe this is due to the distinctive ring structure the PHSs have when $\eta_* = 50$ Mpc, as evidenced by comparing PHS signals distributed in 2D versus in 3D, and by studies testing the CNN on “dot” and “ring” test signals (Appendix C).

After developing the CNN for 90^2 pixel images, we apply it to larger 720^2 pixel maps, sliding 90^2 “templates” in 5-pixel steps across the larger images to generate a probability map. In the probability map, each pixel is evaluated by the network multiple times. As a final step, we filter the probability map, only retaining clusters—groups of positive network outcomes—of a certain size. The benefit of the sliding template search is that it is less sensitive to the exact position of the hotspot within the 90^2 pixel region. Applied in this manner, we find that the CNN can efficiently discern the presence of hotspots, even if the signal temperature is much smaller than the CMB temperature fluctuations. In particular, the CNN can even identify $\mathcal{O}(10)$ number of PHSs on the CMB sky for $g = 1$ and $\eta_* = 50$ Mpc, a signal that has a temperature ≈ 20 times colder than the average CMB temperature fluctuations. Translated into model parameters, for the benchmark models we study using mock CMB maps, we project that a CNN search can set a lower bound on the mass of heavy scalars $M_0/H_I \gtrsim 110\text{--}260$, with the precise value depending on the time of particle production and coupling to the inflaton. These numbers are a significant improvement over the simplistic analysis in Ref. [27] that used a single temperature cut to separate the signal from the background.

Compared to the standard matched filter analysis, the CNN is more versatile in identifying nonrotationally symmetric signals with varying shapes and temperatures that arise in the context of PHSs. We performed a simplified comparison between the CNN and matched filter analysis by considering PHSs with a fixed profile and located on the last scattering surface to show that the match filter analysis can provide comparable signal capture and fake rates to the CNN search for PHSs with $\eta_* = 160$ Mpc and 100 Mpc. For dimmer PHSs ($\eta_* = 50$ Mpc), more analysis is required to separate the signal and background in the

filtered map. We leave a more detailed comparison to the matched filter method with a bank of filters to cover the signal space to future work.

Several future directions remain to be explored. It would be interesting to apply our methodology to actual Planck CMB maps to search for PHSs. In the absence of a detection, we can still set a lower bound on the masses of ultraheavy particles which are otherwise very difficult to discover or constrain. This, however, requires a subtraction of the astrophysical foregrounds and knowing if the CNN can distinguish PHSs from the compact objects in the foreground. Since the distortion of the curvature perturbation from particle production also modifies structure formation at late times, it would also be interesting to see if the current or future large scale structure surveys can identify the resulting signals localized in position space. A neural network like the one used here can learn to incorporate the nonlinear physics of structure formation if trained on suitable simulations. Related to localized PHS signatures, similar types of cosmological signals from topological defects [52] or bubble collisions [29–31] can also arise and these may also be identified by a CNN search. From a more theoretical perspective, it would also be useful to write down a complete inflationary model that incorporates inflaton coupling to heavy fields and leads to particle production as described here. We leave these directions for future work.

ACKNOWLEDGMENTS

We thank Raphael Flauger, Daniel Green, Matthew Johnson, Kin-Wang Ng, Bryan Ostdiek, LianTao Wang, Yiming Zhong for useful conversations. T. K., A. M., and Y. T. are supported by the U.S. National Science Foundation (NSF) Grant No. PHY-2112540. J. K. is supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (Grant No. 2021R1C1C1005076), and in part by the international cooperation program managed by the National Research Foundation of Korea (Grants No. 2022K2A9A2A15000153 and No. FY2022). S. K. is supported in part by the NSF Grant No. PHY-1915314 and the U.S. Department of Energy (DOE) Contract No. DE-AC02-05CH11231. M. M. is supported by the U.S. Department of Energy, Office of Science, under Award No. DE-SC0022342.

APPENDIX A: SENSITIVITY TO THE Λ CDM PARAMETERS

Our analysis uses Λ CDM parameters in Eq. (23) to simulate the CMB. As the Λ CDM parameters come with uncertainties, we should check how sensitive the signal capture rate is to the variation of the parameters. In Table IV, we show the background rejection and signal capture rate using the same trained network for Fig. 8 (left)

TABLE IV. The response of the signal capture and background rejection rates with varying Λ CDM parameters, labeled with the difference to the Λ CDM parameters. The variation of the rates is comparable to the fluctuations in our CNN analysis due to finite sampling and therefore is insignificant. For this test, we used $g = 2$ and $\eta_* = 160$ Mpc for the PHS signal.

	ω_b	ω_{cdm}	$10^9 A_s$	n_s	τ_{re}	Bg rejection	Signal capture
Planck18	0.0224	0.120	2.10	0.966	0.0543	99.8%	74.0%
Case 1		+0.004				99.8%	72.3%
Case 2			+0.07			99.2%	74.1%
Case 3				+0.01		99.6%	69.9%
Case 4	+0.0003					99.8%	73.4%
Case 5					+0.014	99.2%	74.4%
Case 6	+0.0003	-0.004	+0.05	-0.01	-0.014	99.8%	72.4%

with $g = 3$ and $\eta_* = 160$ Mpc but on CMB maps simulated with variations of Λ CDM parameters. As we see, when changing the $\{A_s, \Omega_b, \Omega_{\text{CMB}}, n_s\}$ one by one with twice the 1σ uncertainty reported in [40], the signal capture rate only changes by $\mathcal{O}(\text{few}\%)$, comparable to the variations in our CNN analysis due to finite sampling. The consistent search results show the robustness of the network's ability to identify PHSs against the uncertainty of Λ CDM parameters.

APPENDIX B: PHS CORRECTIONS TO THE CMB POWER SPECTRUM

Here we show the corrections on the CMB power spectrum when the number of PHSs in the full sky saturates the bounds in Table II. We show examples with the coupling $g = 1$ and horizon sizes $\eta_* = 100$ Mpc ($N_{\text{PHS}} = 840$) and 160 Mpc ($N_{\text{PHS}} = 1162$), assuming the centers of all the hotspots are located on the last scattering surface. Notice that the latter

assumption of fixing $\eta_{\text{HS}} = \eta_{\text{rec}}$ makes the average PHS temperature higher compared to the main analysis that allows η_{HS} to vary. However, the assumption simplifies the power spectrum calculation and gives a more conservative result by exaggerating the PHS correction to the power spectrum. We also check results for different g and η_* but, following Table II, with a much smaller N_{PHS} . The corrections to the power spectrum for the other benchmarks are even smaller.

To see how the excesses appear on the power spectrum, we utilize hierarchical equal area isolatitude pixelization, HEALPix [41], based on the C_ℓ^{TT} spectrum computed from the CLASS package using the same Λ CDM parameters in Eq. (23). HEALPix pixelates a sphere in an equal area where the lowest resolution consists of 12 baseline pixels. The resolution is increased by dividing each pixel into four partitions which can be parametrized as $N_{\text{pixels}} = 12N_{\text{side}}^2$ where N_{side} is a power of 2. We choose the resolution parameter $N_{\text{side}} = 2048$. Since the total number of pixels in

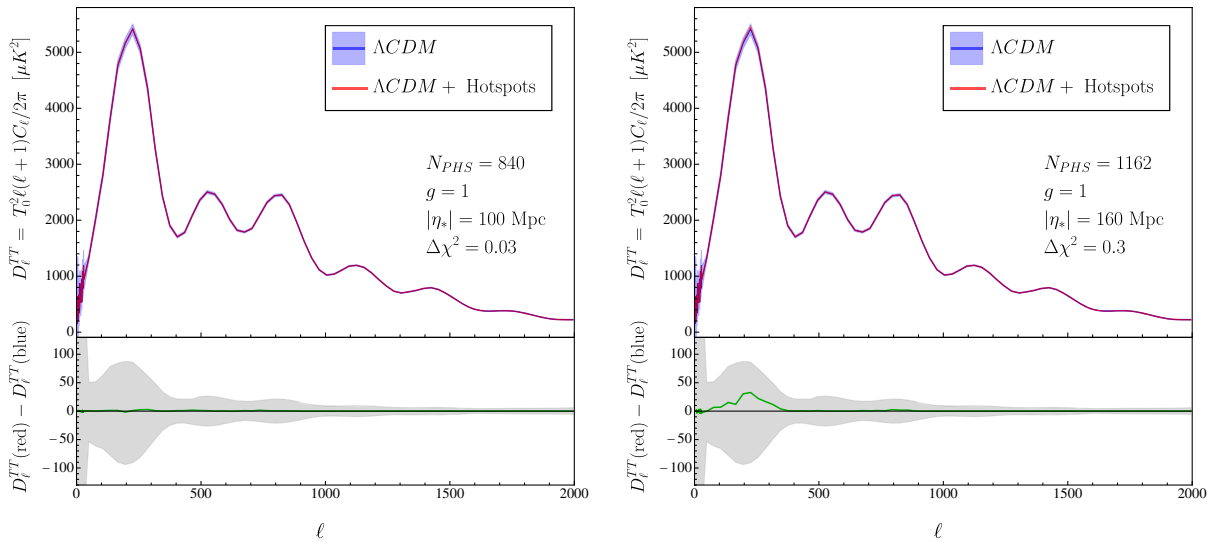


FIG. 13. CMB temperature power spectrum using best fit Λ CDM input parameters in Eq. (23) with (red lines) and without (blue lines) PHS signals implemented on the full sky using a resolution parameter $N_{\text{side}} = 2048$. Here, we assume that all PHS signals are on the last scattering surface. The differences between the two distributions are shown in green lines, and the gray shaded regions denote 1σ uncertainty, taken from the Planck 2018 data.

a sphere characterizes the total number of independent ℓ modes in C_ℓ^{TT} , which is given by $\sum_{\ell=0}^{\ell_{\text{max}}} (2\ell + 1) = (\ell_{\text{max}} + 1)^2$, our benchmark resolution parameter $N_{\text{side}} = 2048$ corresponds to the maximum multipole number $\ell_{\text{max}} \simeq 3500$.

Figure 13 shows $\mathcal{D}_\ell^{\text{TT}}$ spectra for the Λ CDM model (blue) and the Λ CDM + PHS (red) with $\eta_* = 100$ Mpc and $\eta_* = 160$ Mpc. The difference between the red and blue spectra is shown on the lower panel (green), with the 1σ error bar (gray) taken from the *Planck* 2018 result [40]. For both scenarios, the excesses are well below the error bar indicating that the power spectrum analysis will not be able to resolve them. We also show $\Delta\chi^2$ to quantify the

deviations with respect to the Λ CDM spectrum using the same *Planck* 2018 binning intervals in ℓ . The total $\Delta\chi^2$ for both cases is negligible compared to the number of parameters we have.

APPENDIX C: SHAPE ANALYSIS FOR THE $\eta_* = 50$ Mpc SIGNAL

In our earlier results, we found that the CNN's performance for $\eta_* = 50$ Mpc PHSs exceeds the other benchmarks, despite the fact that the hotspots at $\eta_* = 50$ Mpc are much cooler. We surmise that the result is due to the distinct shape of the profile—a rim structure with central peak. As a

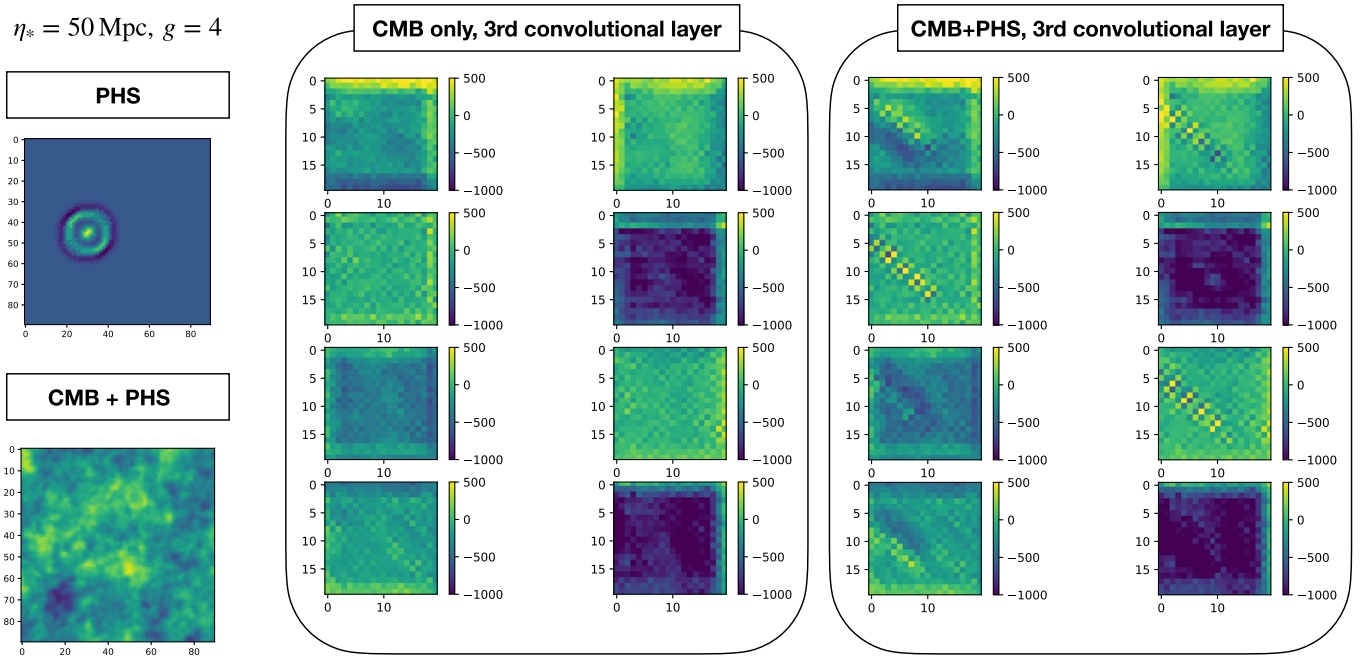


FIG. 14. Similar plots as Fig. 7, but with $\eta_* = 50$ Mpc.

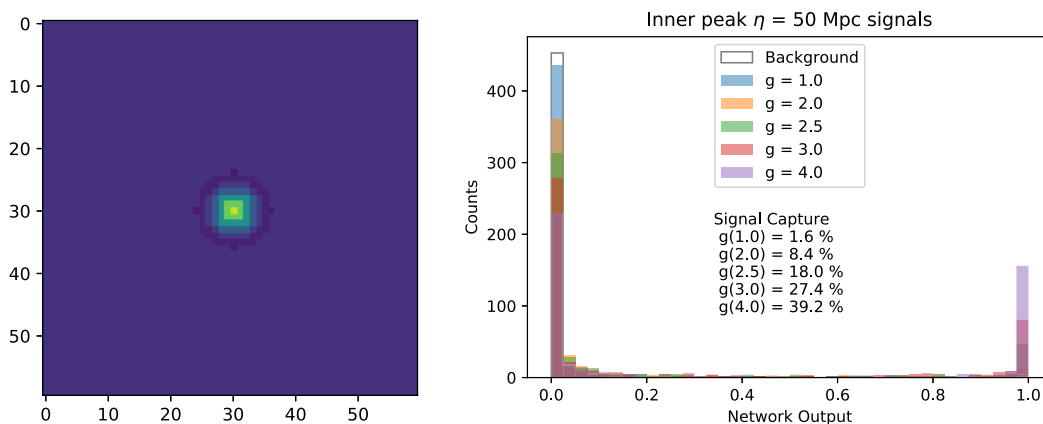


FIG. 15. In the left panel we show the trimmed inner piece of a hotspot signal, while in the right we show the output after 500 CMB + inner hotspot images are run through a network trained on full (untrimmed) $\eta_* = 50$ Mpc, $g = 3$ hotspots.




FIG. 16. In the left panel we show the trimmed outer piece of a hot spot signal, while in the right we show the output after 500 CMB + outer hot spot images are run through the same network as in Fig. 15. The capture rates for the ring are much higher than for the central spot, shown in Fig. 15.

simple test of this hypothesis, we formed a signal set of PHSs decomposed into two separate features, an inner peak and an outer rim. We then ran each piece through a network trained on the complete shape of the $\eta_* = 50$ Mpc spots.

We ran 500 CMB + deconstructed PHS test samples through the network, using a variety of g values but always with both located on the last scattering surface. The results, along with sample images of the deconstructed signals, are shown in Figs. 15 and 16. Comparing the right-hand panels in Figs. 15 and 16, we see that the network is much more efficient at capturing the ring portion, e.g., 88% capture for $g = 3$ compared to 27% for the central spot. From this test we conclude that the ring shape is crucial to the CNN’s performance at low η_* [note that the signal capture for the ring nearly matches the capture rate for the full signal (Fig. 8)]. In Fig. 14, we show the CNN feature maps both with (right) and without (middle) implanted signals for the $\eta_* = 50$ Mpc case. Interpreting the convolution layers’ signal identification is challenging; however, the signal images exhibit more extended profiles compared to the $\eta_* = 160$ Mpc case in Fig. 7, where PHS profiles peak at each of the hot spot centers.

APPENDIX D: ROBUSTNESS OF CNN ANALYSIS WITH DEFECTS IN THE IMAGE

To evaluate the CNN search’s generalizability on real data, we conduct a preliminary test using imperfect simulated CMB maps with PHSs. These simulations incorporate masked regions and randomly distributed bright pixels. To ensure consistency with our analysis, we adopt the true mask utilized in the *Planck* 2018 analysis [53], downscale its resolution to 1024 from the original $N_{\text{side}} = 2048$ setting, and project it onto a flat map. For comparison with the 90^2 -pixel analysis in Sec. IV B, we

randomly sample small mask patches from the flattened map and inject them into sample images containing 90^2 pixels for CNN training and testing purposes.

Each sample image contains masks that collectively cover up to 10% of the image area, and the masks can cut out some PHS signals in the image. Additionally, random pixels are included, with numbers ranging from 0–3, and temperatures varying between 60–130, μK . These temperature values are chosen to be comparable to the central PHS temperatures found in the benchmark models with $g = 3$, which is the coupling for the network training.

The trained CNN discussed in Sec. IV B encounters confusion due to the presence of masked regions and fake point sources. We therefore retrain the network using a fresh set of images that incorporate both masks and bright pixels. For injecting the PHS signal, we assume $\eta_{\text{HS}} = \eta_{\text{rec}}$. A more comprehensive investigation of the effects of these defects, considering variations in PHSs’ line-of-sight distance, will be left for future work.

In Fig. 17, we show the network output for the 1k images with and without injecting the PHS signal. We perform the analysis assuming $\eta_{\text{HS}} = \eta_{\text{rec}}$ for the PHS signals. The obtained background rejection and signal capture rates should be compared to the right panel of Fig. 8. As we can see, the retrained CNN achieved similar background rejection and signal capture rates. For instance, taking $\eta_* = 50$ Mpc as an example, the 2σ upper bounds on the number of PHS increase by approximately a factor of 2 for $g = 1, 2, 3$, mainly due to slightly lower background rejection rate. However, the resulting lower bound on M_0/H_I only drops by up to 2%. This similar sensitivity to the analysis in Sec. IV B confirms the feasibility of adapting the CNN analysis to detect PHSs in real data.

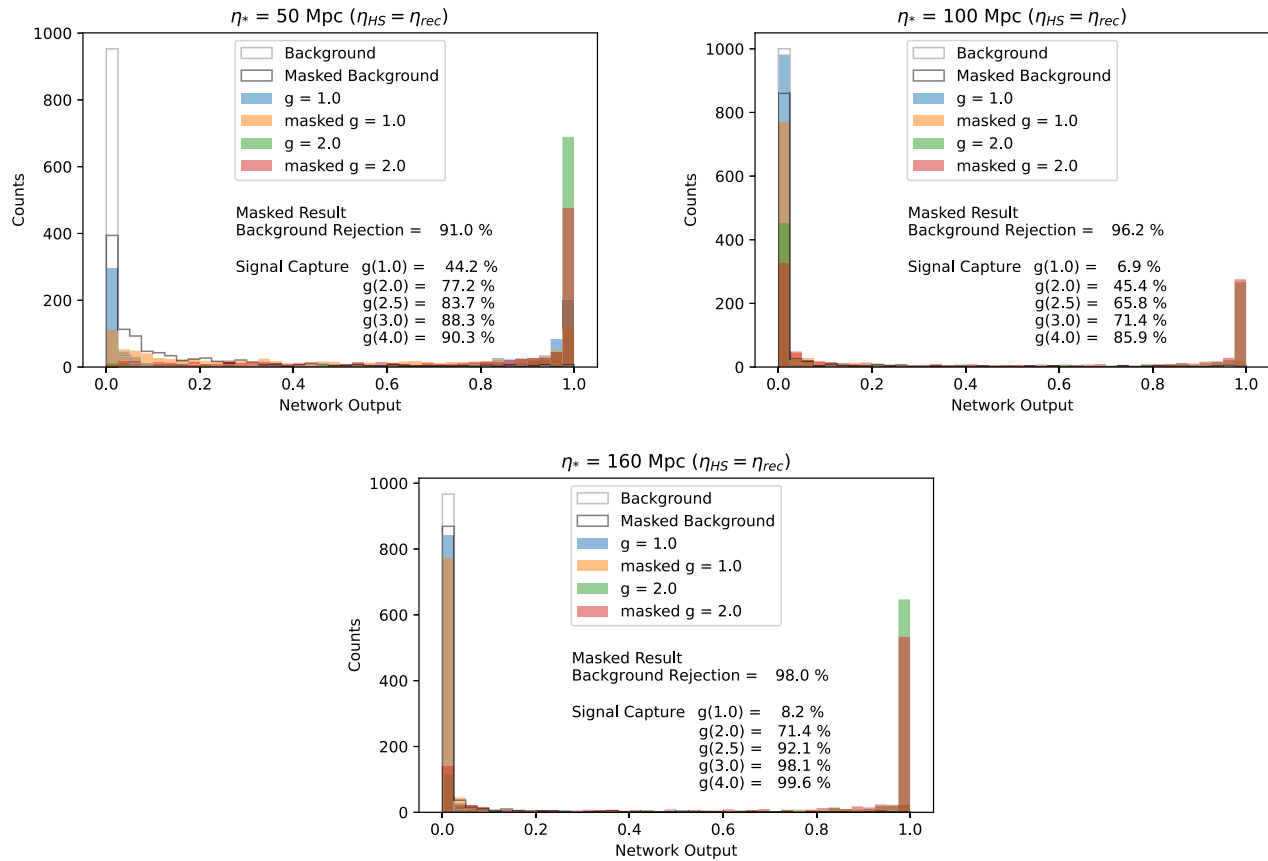


FIG. 17. Training result of CNN with defects.

- [1] A. H. Guth, The inflationary universe: A possible solution to the horizon and flatness problems, *Phys. Rev. D* **23**, 347 (1981).
- [2] A. D. Linde, A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems, *Phys. Lett.* **108B**, 389 (1982).
- [3] A. Albrecht and P. J. Steinhardt, Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking, *Phys. Rev. Lett.* **48**, 1220 (1982).
- [4] Y. Akrami *et al.* (Planck Collaboration), Planck 2018 results. X. Constraints on inflation, *Astron. Astrophys.* **641**, A10 (2020).
- [5] N. D. Birrell and P. C. W. Davies, *Quantum Fields in Curved Space*, Cambridge Monographs on Mathematical Physics (Cambridge University Press, Cambridge, UK, 1984).
- [6] P. Adshead and E. I. Sfakianakis, Fermion production during and after axion inflation, *J. Cosmol. Astropart. Phys.* **11** (2015) 021.
- [7] P. Adshead, L. Pearce, M. Peloso, M. A. Roberts, and L. Sorbo, Phenomenology of fermion production during axion inflation, *J. Cosmol. Astropart. Phys.* **06** (2018) 020.
- [8] X. Chen, Y. Wang, and Z.-Z. Xianyu, Neutrino signatures in primordial non-Gaussianities, *J. High Energy Phys.* **09** (2018) 022.
- [9] L.-T. Wang and Z.-Z. Xianyu, In search of large signals at the cosmological collider, *J. High Energy Phys.* **02** (2020) 044.
- [10] A. Hook, J. Huang, and D. Racco, Searches for other vacua. Part II. A new Higgstory at the cosmological collider, *J. High Energy Phys.* **01** (2020) 105.
- [11] L.-T. Wang and Z.-Z. Xianyu, Gauge boson signals at the cosmological collider, *J. High Energy Phys.* **11** (2020) 082.
- [12] A. Bodas, S. Kumar, and R. Sundrum, The scalar chemical potential in cosmological collider physics, *J. High Energy Phys.* **02** (2021) 079.
- [13] C. M. Sou, X. Tong, and Y. Wang, Chemical-potential-assisted particle production in FRW spacetimes, *J. High Energy Phys.* **06** (2021) 129.

- [14] X. Chen, J. Fan, and L. Li, New inflationary probes of axion dark matter, [arXiv:2303.03406](https://arxiv.org/abs/2303.03406).
- [15] D. J. H. Chung, E. W. Kolb, and A. Riotto, Superheavy dark matter, *Phys. Rev. D* **59**, 023501 (1998).
- [16] D. J. H. Chung, E. W. Kolb, A. Riotto, and I. I. Tkachev, Probing Planckian physics: Resonant production of particles during inflation and features in the primordial power spectrum, *Phys. Rev. D* **62**, 043508 (2000).
- [17] L. Kofman, A. D. Linde, X. Liu, A. Maloney, L. McAllister, and E. Silverstein, Beauty is attractive: Moduli trapping at enhanced symmetry points, *J. High Energy Phys.* **05** (2004) 030.
- [18] A. E. Romano and M. Sasaki, Effects of particle production during inflation, *Phys. Rev. D* **78**, 103522 (2008).
- [19] N. Barnaby, Z. Huang, L. Kofman, and D. Pogosyan, Cosmological fluctuations from infra-red cascading during inflation, *Phys. Rev. D* **80**, 043501 (2009).
- [20] R. Flauger, M. Mirbabayi, L. Senatore, and E. Silverstein, Productive interactions: Heavy particles and non-Gaussianity, *J. Cosmol. Astropart. Phys.* **10** (2017) 058.
- [21] B. Shakya, The Tachyonic Higgs and the inflationary universe, [arXiv:2301.08754](https://arxiv.org/abs/2301.08754).
- [22] X. Chen, Primordial non-Gaussianities from inflation models, *Adv. Astron.* **2010**, 638979 (2010).
- [23] J. Chluba, J. Hamann, and S. P. Patil, Features and new physical scales in primordial observables: Theory and observation, *Int. J. Mod. Phys. D* **24**, 1530023 (2015).
- [24] M. Münchmeyer and K. M. Smith, Higher N-point function data analysis techniques for heavy particle production and WMAP results, *Phys. Rev. D* **100**, 123511 (2019).
- [25] A. Fialkov, N. Itzhaki, and E. D. Kovetz, Cosmological imprints of pre-inflationary particles, *J. Cosmol. Astropart. Phys.* **02** (2010) 004.
- [26] J. Maldacena, A model with cosmological Bell inequalities, *Fortschr. Phys.* **64**, 10 (2016).
- [27] J. H. Kim, S. Kumar, A. Martin, and Y. Tsai, Cosmological particle production and pairwise hotspots on the CMB, *J. High Energy Phys.* **11** (2021) 158.
- [28] B. Hadzhiyska and D. N. Spergel, Measuring the duration of last scattering, *Phys. Rev. D* **99**, 043537 (2019).
- [29] S. Osborne, L. Senatore, and K. Smith, Optimal analysis of azimuthal features in the CMB, *J. Cosmol. Astropart. Phys.* **10** (2013) 001.
- [30] S. Osborne, L. Senatore, and K. Smith, Collisions with other universes: The optimal analysis of the WMAP data, [arXiv:1305.1964](https://arxiv.org/abs/1305.1964).
- [31] J. D. McEwen, S. M. Feeney, M. C. Johnson, and H. V. Peiris, Optimal filters for detecting cosmic bubble collisions, *Phys. Rev. D* **85**, 103502 (2012).
- [32] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XVI. Isotropy and statistics of the CMB, *Astron. Astrophys.* **594**, A16 (2016).
- [33] Y. Akrami *et al.* (Planck Collaboration), Planck 2018 results. VII. Isotropy and statistics of the CMB, *Astron. Astrophys.* **641**, A7 (2020).
- [34] L. Kofman, A. D. Linde, and A. A. Starobinsky, Towards the theory of reheating after inflation, *Phys. Rev. D* **56**, 3258 (1997).
- [35] S. Dodelson and F. Schmidt, *Modern Cosmology* (Elsevier Science, New York, 2020).
- [36] J. Lesgourgues, The cosmic linear anisotropy solving system (CLASS) I: Overview, [arXiv:1104.2932](https://arxiv.org/abs/1104.2932).
- [37] D. Blas, J. Lesgourgues, and T. Tram, The cosmic linear anisotropy solving system (CLASS) II: Approximation schemes, *J. Cosmol. Astropart. Phys.* **07** (2011) 034.
- [38] J. Delabrouille and J. F. Cardoso, Diffuse source separation in CMB observations, in *International Summer School on Data Analysis in Cosmology* (2007), [arXiv:astro-ph/0702198](https://arxiv.org/abs/astro-ph/0702198).
- [39] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2013 results. XII. Diffuse component separation, *Astron. Astrophys.* **571**, A12 (2014).
- [40] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [41] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, HEALPix—A Framework for high resolution discretization, and fast analysis of data distributed on the sphere, *Astrophys. J.* **622**, 759 (2005).
- [42] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).
- [43] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. I. Overview and the cosmological legacy of Planck, *Astron. Astrophys.* **641**, A1 (2020).
- [44] Y. LeCun, P. Haffner, L. Bottou, and Y. Bengio, Object recognition with gradient-based learning, in *Shape, Contour and Grouping in Computer Vision*, Springer, New York, 1999), 319.
- [45] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks, *Adv. Neural Inf. Process. Syst.* **25**, 1097 (2012).
- [46] A. Paszke *et al.*, pyTorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32* (Curran Associates, Inc., Vancouver, Canada, 2019), pp. 8024–8035.
- [47] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [48] X. Glorot, A. Bordes, and Y. Bengio, Deep sparse rectifier neural networks, *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, **15**, 315 (2011), <https://proceedings.mlr.press/v15/glorot11a/glorot11a.pdf>.
- [49] S. Van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Goullart, and T. Yu, scikit-image: Image processing in python, *PeerJ* **2**, e453 (2014).
- [50] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, Asymptotic formulae for likelihood-based tests of new physics, *Eur. Phys. J. C* **71**, 1554 (2011); **73**, 2501(E) (2013).
- [51] N. Craig, F. D’Eramo, P. Draper, S. Thomas, and H. Zhang, The hunt for the rest of the Higgs bosons, *J. High Energy Phys.* **06** (2015) 137.
- [52] H.-T. Cho, K.-W. Ng, and I.-C. Wang, Inflaton fluctuations in the presence of cosmological defects, *J. Cosmol. Astropart. Phys.* **11** (2014) 023.
- [53] NASA/IPAC Infrared Science Archive, Planck map: COM_Mask_CMB-common-Mask-Int_2048_R3.00.