Statistical methods for linking geostatistical maps and transmission models: Application to lymphatic filariasis in East Africa

Panayiota Touloupou[1,2], Renata Retkute[3], T. Déirdre
Hollingsworth[4], and Simon E. F. Spencer[1,2]

[1]Department of Statistics, University of Warwick, Coventry, UK
[2]Zeeman Institute, University of Warwick, Coventry, UK
[3]Department of Plant Sciences, University of Cambridge, Cambridge, UK
[4]Big Data Institute, Li Ka Shing Centre for Health Information and Discovery,
Nuffield Department of Medicine, University of Oxford, Oxford, UK

October 26, 2020

Abstract

Infectious diseases remain one of the major causes of human mortality and suffering. Mathematical models have been established as an important tool for capturing the features that drive the spread of the disease, predicting the progression of an epidemic and hence guiding the development of strategies to control it. Another important area of epidemiological interest is the development of geostatistical methods for the analysis of data from spatially referenced prevalence surveys. Maps of prevalence are useful, not only for enabling a more precise disease risk stratification, but also for guiding the planning of more reliable spatial control programmes by identifying affected areas. Despite the methodological advances that have been made in each area independently, efforts to link transmission models and geostatistical maps have been limited. Motivated by this fact, we developed a Bayesian approach that combines fine-scale geostatistical maps of disease prevalence with transmission models to provide quantitative,

1

spatially-explicit projections of the current and future impact of control programs against a disease. These estimates can then be used at a local level to identify the effectiveness of suggested intervention schemes and allow investigation of alternative strategies. The methodology has been applied to lymphatic filariasis in East Africa to provide estimates of the impact of different intervention strategies against the disease.

Keywords— Bayesian methods; Fine-scale spatial predictions; Linking maps with models; Lymphatic filariasis; Projections; Uncertainty.

1   Introduction

Geostatistical modelling is increasingly used in epidemiology to combine surveys from multiple locations into a detailed model of local prevalence or incidence (Hay et al. 2009, Stensgaard et al. 2011, Moraga et al. 2015, O'Hanlon et al. 2016; Giorgi et al. 2018). Maps of disease distribution can be used, for example, to plan the development of national scale control strategies by informing policy makers where intervention efforts should be focused (Tatem et al. 2010, Slater and Michael 2013). Several examples from the literature have shown that spatial heterogeneity is an important epidemiological factor in many diseases (Sturrock et al. 2010, Pullan et al. 2012, for example), however predictions of future cases are frequently performed on aggregated data, risking the ecological fallacy (Wakefield and Lyons 2010).

Over the past decades, mathematical models have also been established as an important tool for evaluating the effect of different control strategies by predicting the progression of the disease (Ferguson et al. 2005, Tildesley et al. 2009, Stolk et al. 2018a, Hollingsworth 2018). However, when mathematical modelling is used to evaluate potential intervention strategies, spatial heterogeneity is also frequently ignored (Heesterbeek et al. 2015). Some notable exceptions are the papers by Gibson (1997), Keeling et al. (2001) and Deardon et al. (2010) who considered a spatial model, where the transmission probabilities depend on distances between individuals. In this paper we develop a novel method for taking the output from a geostatistical model and projecting the epidemic dynamics forward in time at the pixel level, under a range of potential intervention strategies, in a computationally efficient way. An important feature of our approach is the ability to capture several sources of uncertainty.

2

There are only a limited number of studies linking transmission models and geostatistical maps in a way that can dynamically inform policy at a local level. The African Program for Onchocerciasis Control was one of the first groups to develop and apply this approach (for example Plaisier et al. (1991), Alley et al. (1994)). Kriging was used to extrapolate between survey points and then transmission models were used to project the likely impact of intervention programs. These mapped projections had been extremely useful in informing policy planning over many years and have recently been updated (Tekle et al. 2016). The power of this type of approach to inform policy has been illustrated most notably by Bhatt et al. (2015) in the analysis of the key drivers of successes in malaria interventions over the last 15 years. An important challenge, addressed by our approach, is to appropriately estimate and communicating the projections with their uncertainty. In particular, our method addresses and quantifies a broad range of uncertainties, including uncertainty in the spatial variation in prevalence, transmission parameters, demographics, interventions and even model structure, and propagates them into the uncertainty in future predictions.

Our methodology has many parallels with exact versions of ABC (Wilkinson 2013, Beaumont et al. 2002), in which simulations from the model are weighted (or accepted) according to their likelihood of producing the observed data. However, in our framework a likelihood is only available at the survey points, and so instead we weight the simulations by the posterior distribution from a geostatistical model that interpolates between surveys to give a prevalence distribution at each location. To achieve this weighting, we must change the measure of the simulated prevalences from the one induced by the prior on the transmission model parameters, to the posterior distribution from the geostatistical model using the Radon-Nikodym derivative (Billingsley 1995). However, since this is not available in analytical form, we propose an empirical alternative similar to Goldie and Maller (1999) and references contained within.

The paper is organised as follows. In Section 2 we describe the statistical methodology for combining geo-statistical mapping and transmission modelling, and illustrate its key features with a toy example in Section 3. The proposed method is applied in Section 4 to investigate the impact of intervention programs for lymphatic filariasis in seven countries in Africa. Finally, we conclude with a discussion on limitations of the current method and possible extensions for further research in Section 5.

2   Methods

We develop a Bayesian methodology that captures uncertainty from multiple sources and can be readily applied to different transmission models and intervention strategies to give a distribution of projections across space. The starting point for our analysis is the output from a geostatistical model of disease prevalence, capturing the uncertainty in the spatial distribution of infection. A number of recent studies have adopted a predictive framework known as model-based geostatistics (Diggle and Ribeiro 2007) for the production of prevalence maps, often employing Bayesian inference for spatial prediction and robust characterisation of uncertainty surrounding those predictions. In particular we assume that the output consists of M Monte Carlo samples from the posterior distribution of the geostatistical model. Although we assume that the spatial distribution represents the pre-control prevalence here, our methodology can easily be generalised beyond this example. In addition, we assume that other geographical information is available for each pixel (with associated uncertainty), such as population size and other demographic data that can be used as an input to the transmission model.

   Our methodology consists of 3 steps. First, we generate a large number of simulations from the transmission model, with sufficient variability to capture all of the endemic prevalences observed in the samples from the geostatistical model. Second, for each spatial location in the map we reweight the simulations according to how similar they are to the observed prevalence and other spatial information, such as population data. Finally, we simulate the transmission model further forward in time, possibly under some intervention strategy, and apply the weights to obtain the spatial distribution of the projections. A graphical representation of the method can be found in Figure 1.

2.1   Step 1: Simulating from the transmission model

For each pixel on the map, we assign an informative prior on the model parameters, $\pi_i(\boldsymbol{\theta})$ say for pixel i, representing the uncertainty in our beliefs about the parameters of the transmission model at that location. Next, we define a single proposal distribution over the parameter space, $q(\boldsymbol{\theta})$, capable of producing simulated prevalence levels spanning the values observed in the geostatistical mapping. We then draw J parameter vectors $(\boldsymbol{\theta}_j)$ from the proposal, and for each one we run the model forward in time until it reaches
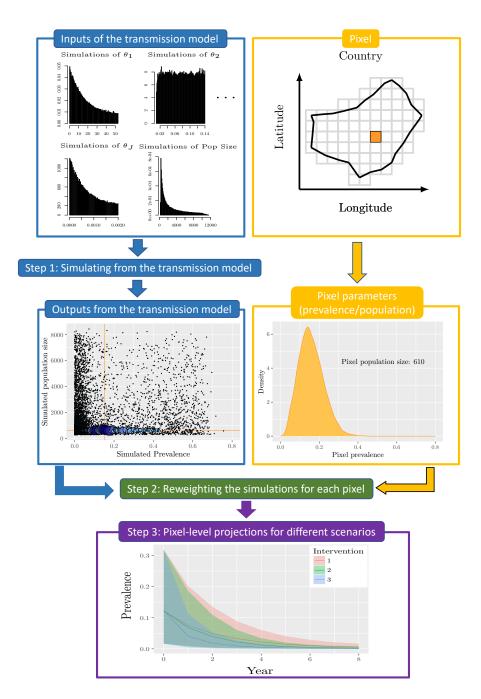
Figure 1: Methodology for generating mapping results. Using pre-run model simulations (top), we reweight the simulations for each pixel based on the prevalence and population information (middle – red lines represent the median of observed data). Finally, the weights are used to evaluate the impact of different intervention strategies (bottom).

pre-control equilibrium. Denote the resulting prevalence levels by p_j, for $j = 1, \ldots J$. Finally, we calculate an initial $I \times J$ matrix of weights for the I pixels and J simulations according to the usual importance sampling formula, namely $w_{ij}^{(1)} = \pi_i(\boldsymbol{\theta}_j)/q(\boldsymbol{\theta}_j)$.

The proposal distribution $q(\boldsymbol{\theta})$ might be uniform over the parameter space in low dimensions, or for higher dimensions it could be developed from pilot simulations, where parameter vectors are sampled uniformly from the support of the priors and for each parameter vector the equilibrium prevalence is simulated from the transmission model. The importance proposal can then be constructed on the parameter space to give more weight to frequently observed prevalence values, and zero weight to implausible prevalence values (e.g. prevalences larger than the maximum observed in the geostatistical model). The efficiency of the proposal can be improved iteratively using adaptive importance sampling techniques (Cornuet et al. 2012).

2.2 Step 2: Reweighting the simulations to match pixel prevalence distributions

For each pixel the same simulations are reweighted to match the prevalence distribution of that pixel. This prevents unnecessary replication of simulations for pixels that are broadly similar and means that the number of simulations need not increase as the number of pixels increases. More specifically, for pixel $i = 1, \ldots, I$ and simulation $j = 1, \ldots, J$ the new weight is given by:

$$w_{ij}^{(2)} \propto \frac{f(p_j|\boldsymbol{d}_i)}{g(p_j|\boldsymbol{w}_i^{(1)})} w_{ij}^{(1)} \tag{1}$$

where $\boldsymbol{d}_i = (d_{i1}, \ldots, d_{iM})$ is the M dimensional vector of posterior samples of prevalences in pixel i and $\boldsymbol{w}_i^{(1)} = (w_{i1}^{(1)}, \ldots, w_{iJ}^{(1)})$. The function f represents the probability of having prevalence p_j under the geostatistical model and g represents the probability of simulating prevalence p_j from the model with parameter vector drawn from the prior. The ratio f/g therefore represents the usual change of measure formula (Radon-Nikodym derivative). However, since neither of these probability densities are likely to be available in closed form, we use an empirical approximation given by the amount of probability

density within $\delta/2$ of p_j:

$$f(p_j|\boldsymbol{d}_i) = \frac{1}{\delta M} \sum_{m=1}^{M} \mathbb{1}_{\{p_j - \delta/2 \le d_{im} \le p_j + \delta/2\}},$$

$$g(p_j|\boldsymbol{w}_i^{(1)}) = \frac{\sum_{k=1}^{J} w_{ik}^{(1)} \mathbb{1}_{\{p_j - \delta/2 \le p_k \le p_j + \delta/2\}}}{\delta \sum_{k=1}^{J} w_{ik}^{(1)}}.$$

Note that as long as $q(\boldsymbol{\theta}) > 0$ implies $\pi_i(\boldsymbol{\theta}) > 0$, then $w_{ij}^{(1)} > 0$ for all j and hence $g(p_j|\boldsymbol{w}_i^{(1)}) > 0$ for all j. The bin width δ controls the trade-off between effective sample size and the fidelity of the distribution of the simulated prevalences to the geostatistical posterior distribution, and should be set as small as possible, whilst producing a reasonable effective sample size, defined as $\left(\sum_{j=1}^{J} w_{ij}^{(2)}\right)^2 \Big/ \sum_{j=1}^{J} \left(w_{ij}^{(2)}\right)^2$. Finally, the weights from Equation (1) are normalised to give the posterior probabilities (according to the geostatistical model) that simulation j is appropriate for pixel i.

2.3 Step 3: Running the simulations forward

The simulations are run forward in time under a given intervention strategy. For each pixel the projected outcomes are weighted according to the normalised weights $\boldsymbol{w}_i^{(2)} = (w_{i1}^{(2)}, \ldots, w_{iJ}^{(2)})$ produced in Step 2. Step 3 is repeated for each intervention strategy under consideration.

2.4 Lemma on the change of measure

In this section we introduce a lemma that generates the reweighting formula in Equation (1). The lemma is proved in Section A.1 of the Appendix.

Lemma 1. *Let $p : \mathbb{R}^d \to [0, 1]$ denote a deterministic model that produces a prevalence $p(\boldsymbol{\theta})$ from a vector of parameters $\boldsymbol{\theta}$. Let $\pi(\boldsymbol{\theta})$ be a prior distribution over the parameters that induces a prior distribution over prevalences, which we denote by $g(p)$. Suppose that there exists a differentiable and invertible function $\boldsymbol{\phi} : \mathbb{R}^d \to \mathbb{R}^d$ that admits the prevalence as its first argument, ie. $\boldsymbol{\phi}(\boldsymbol{\theta}) = (p(\boldsymbol{\theta}), \boldsymbol{q}(\boldsymbol{\theta}))$ for some $\boldsymbol{q}(\boldsymbol{\theta})$. Finally suppose that we wish to change the probability measure over prevalences from g to another measure f that is absolutely continuous with respect to g. Then the resulting measure over the parameter space is given by $h(\boldsymbol{\theta}) = \frac{f(p(\boldsymbol{\theta}))}{g(p(\boldsymbol{\theta}))} \pi(\boldsymbol{\theta})$.*

Notes:

1. The same approach can be applied for stochastic transmission models as long as the model is defined on a separate probability space $(\Omega, \mathcal{F}, P)$ to the prior. For stochastic models we fix $\omega \in \Omega$ and consider the transmission model as a deterministic map $\boldsymbol{\phi}(\boldsymbol{\theta}, \omega)$, applying the Lemma and then integrating over Ω.

2. The condition that f must be absolutely continuous with respect to g means that whenever $g(p) = 0$ then we must also have $f(p) = 0$. In other words, when the prior probability of a prevalence is zero then the map measure of prevalence must also be zero. This has important implications for the implementation of our method, discussed further in Section A.2.

2.5 Alternative empirical Radon-Nikodym derivatives

In Step 2 of our algorithm (described in Section 2.2) we proposed an empirical estimate of the Radon-Nikodym derivative f/g based on using the prevalences within $\delta/2$ of p_j. Clearly there are many possible alternative estimates that could be used and there are two in particular that are worthy of further discussion. The first is based on histograms and the second is based on minimising a discrepancy measure.

2.5.1 Histogram-based empirical Radon-Nikodym derivative

If we consider a fixed partition of the prevalence space into bins (as if we were constructing a histogram) then it is straightforward to calculate the Radon-Nikodym derivative f/g for each bin as the proportion of posterior samples in the bin divided by the proportion of the weight belonging to simulations that fall in the bin. More precisely, given a finite set of disjoint intervals with union $[0, 1]$ then if prevalence p_j falls in interval $\mathcal{I}(p_j)$ we have that

$$f(p_j|\boldsymbol{d}_i) = \frac{1}{M|\mathcal{I}(p_j)|} \sum_{m=1}^{M} \mathbb{1}_{\{d_{im} \in \ \mathcal{I}(p_j)\}},$$

$$g(p_j|\boldsymbol{w}_i^{(1)}) = \frac{\sum_{k=1}^{J} w_{ik}^{(1)} \mathbb{1}_{\{p_k \in \ \mathcal{I}(p_j)\}}}{|\mathcal{I}(p_j)| \sum_{k=1}^{J} w_{ik}^{(1)}},$$

where $|\mathcal{I}(p_j)|$ is the length of the interval containing p_j.

The main advantage of this estimate is computational – since all of the simulations in the same interval have the same ratio (for a given pixel) then instead of having to calculate J ratios we need only calculate one per interval. A secondary advantage is that the weighted histogram of the simulation prevalences will be identical to the histogram of the posterior prevalence distribution. However, the relative weightings within each bin are unchanged and so a different choice of bins will reveal that the two distributions are different.

2.5.2 Discrepancy-based empirical Radon-Nikodym derivative

A second alternative empirical Radon-Nikodym derivative can be defined to minimise the difference between the empirical cumulative distribution functions (cdfs) of the posterior prevalences and the weighted simulated prevalences. Let $F(x|\boldsymbol{d}_i) = \frac{1}{M} \sum_{m=1}^{M} \mathbb{1}_{\{d_{im} \leq x\}}$ be the empirical cdf of the map prevalence distribution for pixel i and $H(x|\boldsymbol{w}_i^{(2)}) = \sum_{j=1}^{J} w_{ij}^{(2)} \mathbb{1}_{\{p_j \leq x\}}$ be the empirical cdf of the final weighted distribution of simulated prevalences, then we can choose $\boldsymbol{w}_i^{(2)}$ to minimise some distance $||F(\cdot|\boldsymbol{d}_i) - H(\cdot|\boldsymbol{w}_i^{(2)})||$. For example, we may wish to minimise $\int_0^1 |F(x|\boldsymbol{d}_i) - H(x|\boldsymbol{w}_i^{(2)})| \, \mathrm{d}x$ or $\int_0^1 \left(F(x|\boldsymbol{d}_i) - H(x|\boldsymbol{w}_i^{(2)})\right)^2 \, \mathrm{d}x$. In this paper we have focussed on the latter of these, for details of the calculation we refer the reader to the Appendix A.3.

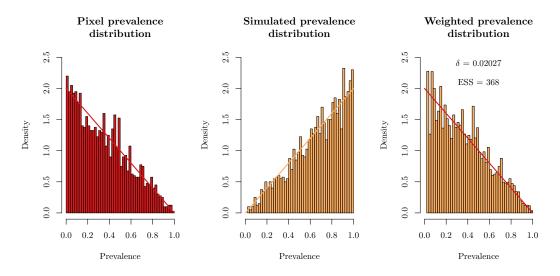## 3 Simulation studies: A toy example

In this section we provide a toy example to assess the performance of the proposed method under different settings. Particular focus was given on how the method was affected by the value of δ, by the choice of the proposal distribution of the parameters and the empirical estimate of the Radon-Nikodym derivative. A full description of the analysis can be found in Appendix B and here we summarize the key results.

Suppose that the prior distribution is $\pi(\theta_1, \theta_2) = 2$ if $0 < \theta_2 < \theta_1 < 1$ and zero otherwise. Plots illustrating this prior are given in Appendix Figure B.1. For simplicity, assume that the transmission model has equilibrium prevalence given by $p(\theta_1, \theta_2) = \theta_1$ so that the induced prior over prevalences

is the marginal for θ_1, ie. $g(p) = 2p$ for $0 < p < 1$, which is a Beta(2,1) distribution. Further, suppose that we are given $M = 2000$ samples from a pixel with prevalence measure $f(p) = 2(1 - p)$ for $0 < p < 1$, representing a Beta(1,2) distribution. This challenging example allows us to assess how the methodology performs when there are few simulations with low weights in the area of high posterior probability close to $p = 0$.
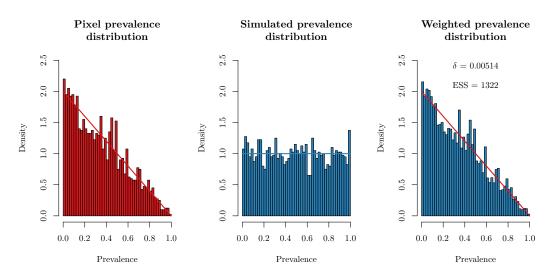
Simulations were conducted to investigate the accuracy and efficiency of the proposed method under different settings, where the observed pixel and simulated prevalence data are obtained from the toy model. Figure 2 shows how $J = 2000$ simulations from a proposal (centre histogram) can be reweighted (right histogram) to resemble the pixel prevalence distribution (left histogram). In Figure 2(a) the proposal is from the prior, whilst in Figure 2(b) the proposal is $U(0, 1)$. The improvement due to the proposal distribution having good support in all areas of the posterior distribution was demonstrated by the substantial increase in effective sample size (ESS; from 368 to 1322), despite a much smaller value of δ.

Figure 3 illustrates how the performance changes as δ is increased. The left figure shows the distance (given by integrated squared difference) between the empirical cumulative distribution functions of the weighted simulations and the samples from the pixel posterior; and the right figure shows the effective sample sizes. The corresponding results from the discrepancy based empirical Radon-Nikodym derivative (see Section 2.5.2) are shown as horizontal dashed lines and provide the minimum distance possible between the cdfs. The results show that smaller values of δ reproduced the empirical cdf more accurately, unless δ was so small that very few simulations were included in each estimate of the density g (the denominator in Equation (1)). After some experimentation (see Appendix B.1.1) we chose to set δ to be the smallest value for which at least three simulations were included in each estimate of g. These values are illustrated by vertical lines on Figure 3, and can be seen to come close to achieving the minimum possible squared distance between the empirical cdfs, but with larger ESS.

In our simulations we have evaluated the performance of the method for the distance-based empirical Radon-Nikodym derivative, which is based on using the prevalences within a certain distance given by $\delta/2$ of p_j. We also investigated alternative derivatives, discussed in Section 2.5, and the results are summarised in Table B.1 of the Appendix. Overall, we observed that the discrepancy-based derivative provides the best possible distance between the two cdfs, but at a cost of a lower ESS in all the scenarios considered. When

(a) The proposal yields a simulated prevalence distribution equal to its marginal prior, Beta(2,1).

(b) The proposal yields a uniform simulated prevalence distribution.

Figure 2: The estimated weighted prevalence distribution for the suggested value of δ (right panel) is compared to the true pixel prevalence distribution (left panel), under different proposal distributions for the prevalence (middle panel): (a) Beta(2,1) and (b) $U(0,1)$. The target densities are also shown on each panel.
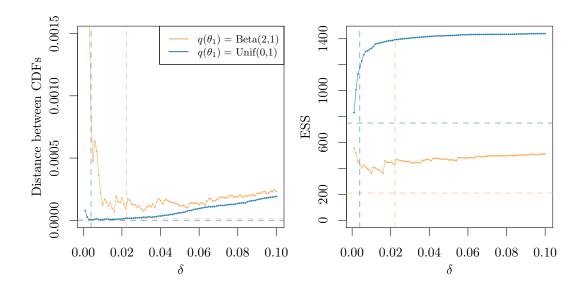
11

Figure 3: Distance between the two cumulative distribution functions (cdfs) (left panel) and effective sample size (ESS) (right panel) obtained under different values of δ and choice of proposal distribution for parameter θ_1, for one randomly selected simulated dataset. Orange solid line represents a prevalence proposal distribution equal to the marginal prior, i.e. Beta(2,1), whereas the blue line corresponds to a $U(0,1)$ proposal. In both cases, the pixel prevalences were drawn from a Beta(1,2) distribution. Dashed vertical lines represent the suggested value of δ for each scenario considered. Dashed horizontal lines correspond to the minimum possible distance (left panel) and its associated ESS (right panel).

the proposal was uniform and had simulations in all areas of the posterior distribution then the histogram-based derivative performed better than the distance-based derivative both in terms of accuracy and ESS. However, the situation was reversed when there were areas in the proposal with few simulations. In that scenario, the distance-based derivative was found to have lower integrated squared distance and higher ESS compared to the histogram-based derivative. Therefore, we used the distance-based empirical Radon-Nikodym derivative in our applications, since it was more robust to weaknesses in the proposal.

4 Application to lymphatic filariasis data

In this section, we apply the proposed approach for the analysis of real data for lymphatic filariasis (LF) in East Africa. LF is caused by a mosquito-borne macro-parasite, which was historically endemic in many tropical countries, with over a billion people at risk of infection, and millions affected by the disease suffer from disability, stigma and associated social and economic consequences (Ramaiah and Ottesen 2014). LF is one of the neglected tropical diseases (NTDs) targeted for elimination as a public health problem by 2020 (WHO 2012), with new guidelines currently being developed for 2030. Global efforts to eliminate LF as a public health problem, through the use of mass drug administration (MDA) of treatments with an excellent safety record, have reduced prevalence to low levels in many settings (Ramaiah and Ottesen 2014). While many countries have successfully scaled-up their programs, there remain a number of questions on how best to scale up treatment to assist priority countries in optimising interventions to accelerate elimination. Therefore, there is an urgent need to provide detailed estimates of the impact of current and future control programs for donor and policy planning.

For LF, the intervention strategy for most of Africa is to have yearly MDA at 65% coverage for 5 years, followed by an assessment of transmission and, if necessary, further rounds of treatment. In areas where MDA has not yet started, alternative strategies may be required to meet the WHO target, i.e. the prevalence being less than 1% (WHO 2012) as soon as possible. Enhanced strategies include MDA at high coverage or twice-yearly treatment (Stolk et al. 2018b). By bringing together statistical mapping and transmission modelling, we aim to provide high-resolution quantification of the likely impact of control programs and predictions on both future impact and demand for interventions, allowing policy makers to more effectively target available resources.

4.1 The mathematical model of LF transmission dynamics

In this section we describe the mathematical model of lymphatic filariasis transmission, TRANSFIL (Irvine et al. 2015), that is used throughout the paper. TRANSFIL is an individual-based model of LF infection in human populations, with each host having their own adult worm and microfilariae (mf) burden, as well as mosquito bite risk and treatment history. A full

description of the model is provided in Irvine et al. (2015) and in Section C of the Appendix, so here we provide only a brief overview and the updated aspects of it.

Each human is assumed to have their own burden of male and female worms denoted by W_i^m and W_i^f, respectively. The times at which human i acquires female and male adult worms are given by two inhomogenous Poisson processes, both with rate:

$$\tfrac{1}{2}\lambda b_i(V/H)\psi_1\psi_2 s_2 h(a),$$

where λ is the number of bites per mosquito, V/H is the ratio of vectors to hosts, ψ_1 is the probability that a third-stage larvae (L3, the infectious stage) leaves the host during a bite, ψ_2 is the probability that the L3 enters the host, s_2 is the proportion of L3 that develop into adult worms within the host and $h(a)$ is the biting rate for a human with age a. Both male and female worms are introduced to a human according to a bite risk b_i drawn from a gamma distribution with mean 1 and shape parameter k. Thus, the degree of parasite aggregation amongst humans can be quantified by this shape parameter. Finally, we assume that each worm has a constant death rate μ.

Microfilariae concentration in the peripheral blood, denoted by M_i, is also modelled for each individual according to the following equation:

$$\frac{\mathrm{d}M_i}{\mathrm{d}t} = \alpha W_i^f \mathbb{1}_{\{W_i^m>0\}} - \gamma M_i,$$

with α being the production rate of mf per worm, γ the constant death rate of mf and the indicator function $\mathbb{1}_{\{W_i^m>0\}}$ is one if there are male worms and zero if not. Larvae development occurs when mf entered the mosquito during a blood meal from an infected host. Different functional forms have been found to describe the relationship between the number of mf ingested and the number that develop within the mosquito. For *Anopheles*, which is the genus of the most dominant vector species in East Africa, this relationship is expressed as:

$$L(m) = \kappa_{s2}\left(1 - e^{-r_2 m/\kappa_{s2}}\right)^2,$$

where m is the concentration of mf per 20μL taken during a blood meal and r, κ_s denote the saturation values related to the uptake function as detailed in Gambhir and Michael (2008). The equilibrium value for L3 in a mosquito

14

is given by:

$$L^* = \frac{\lambda g \tilde{L}}{\sigma + \lambda \psi_1},$$

where λ is the number of bites per mosquito, g is the proportion of mosquitoes which pick up infection when biting an infected host, σ is the death rate of mosquitoes and $\tilde{L}$ is the average number of larvae per mosquito.

Each human begins life with zero infection and a bite-rate of exposure b_i. The human death rate is denoted by τ and is assumed to be constant throughout an individual's lifetime with a cut-off at age 100. When an individual dies another one is born in order to keep the population size constant.

During an intervention campaign, the impact of MDA is simulated for an individual by reducing their mf concentration and their male and female worm burden according to the estimated drug efficacies from the literature (Ismail et al. 1998, Michael et al. 2004). In addition, there is a period after MDA during which the production of mf for that individual is diminished. Furthermore, the individuals' compliance after multiple rounds of treatment is modelled based on the paper by Griffin et al. (2010), where the authors model the probability of an individual making the same decision as in the previous round of treatment.

Finally, we extended the model to include a very low rate of importation of infection from outside the population being modelled, otherwise the equilibrium distribution (steady state), that is used as the starting point of the simulations, is just the degenerate distribution where no-one is infected. The interventions reduce the prevalence over time, and so we reduce the importation rate after intervention in proportion to the reduction in prevalence seen in pilot simulations. Lists of the model parameters are provided in Tables C.1 and C.2 of the Appendix.

4.2 Implementation details

The starting point for our analysis was the spatial map (pixel scale 5×5 km) providing the predicted distribution of the LF prevalence based on mf data, generated through a Bayesian geostatistical modelling approach described by Moraga et al. (2015). The top panel of Figure 4 shows the median of the posterior distribution of the prevalence obtained at each pixel, along with estimates of lower (2.5%) and upper (97.5%) quantiles. In particular,

15

we analysed the following seven African countries: Ethiopia, Sudan, South Sudan, Eritrea, Kenya, Tanzania and Uganda. We linked each pixel to the corresponding population estimates obtained from the Gridded Population of the World (Worldpop 2010), which provides the estimated number of people in each pixel. We avoided handling pixels with either very small or very large populations as the transmission model was not thought to be appropriate in these environments (Irvine et al. 2015, Smith et al. 2017). More specifically, for small populations we pooled pixels with less than 300 people together, ensuring that the merged pixels belong to the same country and that the groups contain as few pixels as possible. We excluded pixels with population estimates over 10 000 from the analysis; resulting in 1.7% of the pixels being excluded.

The stochastic model of LF transmission TRANSFIL was used to investigate and compare the impact of different control strategies. In order to simulate the entire range of observed baseline mf prevalence levels, with values up to 95%, we assumed that four parameters of the mathematical model were spatially varying: the population size, the vector to host ratio, the aggregation parameter of individual exposure to mosquitoes and the importation rate, using prior distributions informed from data, pilot simulations and previous analyses. We assumed that the parameter prior was the same in each spatial location (discussed in more details in Section C.2 of the Appendix) except for the population size, which was assumed to be a log normal distribution, ie $\log(n) \sim \mathcal{N}\left(\log(N_i), \sigma^2\right)$, where N_i is the reported population of pixel i (adjusted population from Worldpop 2010) and σ is the sample standard deviation of the log population estimates available in WorldPop.

The proposal density of the population sizes, $q(n)$, was designed so that each simulation contributed an equal amount to the effective sample size of a set of pixels with populations $\{260, 261, \ldots, 10\,000\}$. This was achieved by calculating the effective sample size of an initial proposal, namely, $q_0(n) \propto 1$. The remaining populations sizes ($10\,001$–$11\,550$) were taken to decrease linearly from $q(10\,000)$ to zero. Since the uncertainty in the log-normal prior is much greater for large populations, fewer simulations are needed in these regions. The final proposal was obtained from 10 iterations of $q_i(n) \propto q_{i-1}(n)/\mathrm{ESS}_{i-1}(n)$, where $\mathrm{ESS}_i(n)$ is the effective sample size of a simulation with n individuals from the proposal q_i (see Figure C.1(a) of the Appendix).

A significant merit of our approach is that it can be easily applied in parallel which can be utilised to speed up implementation, especially in applications involving a large number of pixels. This is because we are treating

each pixel independently and therefore the computation of the weights can be undertaken in parallel. In our application, the computation time of this step was approximately 8 hours using a 112 core computer cluster.

4.3   Results

In this section, the Bayesian approach presented in Section 2 was applied on the LF data. Firstly, we assessed the accuracy of the method, defined as the ability of the transmission model weighted simulations to reproduce the pre-control (baseline) geostatistical map, by comparing the observed and the estimated distribution of the baseline (equilibrium) mf prevalences at each pixel. Figure 4 illustrates the median map (with 2.5 and 97.5 percentiles), along with the corresponding maps of the observed data. Overall, the results show that the maps are almost identical, indicating that the method is able to reproduce the distribution of the observed baseline prevalence in each pixel. In addition, in the left panel of Figure D.1 of the Appendix we compared the estimated number of people per pixel with the observed value, which were in close agreement indicating that the proposed method accurately reproduced the number of people in each pixel. In the right panel of Figure D.1 of the Appendix, we examined the ESS per pixel, which represents the effective number of simulations per pixel and is a measure of how well the method performs. We observed that the pixels with high prevalence (which may require a change of intervention strategy) have high ESS.

Secondly, the methodology was applied to evaluate the impact of different intervention programs for LF in East Africa. In particular, four treatment scenarios were simulated: no interventions; the standard 65% coverage annual MDA (aMDA); 80% coverage aMDA; or biannual MDA (bMDA) at 65% coverage, in order to investigate how these affect the probability of elimination after 5 years (Figure 5). Adopting a prevalence of less than 1% as the threshold set by WHO as a global target for determining LF transmission elimination, our analysis predicted that the recommended strategy of 5 rounds of aMDA at 65% is not enough for eliminating the disease in all pixels, with probability of elimination above 90% only for 13% of the pixels (see also Figure 6). Moreover, when more intensive treatments were implemented, i.e. more frequent MDA or higher coverage, the probability of elimination significantly increased compared to aMDA programme at 65%. In particular, bMDA at 65% coverage was the most effective of all strategies considered and was able to achieve elimination in 88% of the pixels, with at
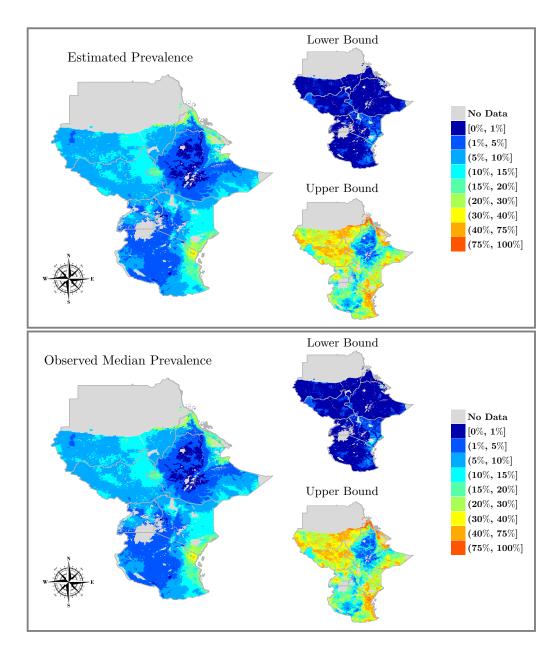
Figure 4: Accuracy assessment of our method. The true distribution of LF prevalence (lower panel) is compared to the estimated prevalence distribution (upper panel) using our proposed methodology predicted at 5×5 km resolution. Point estimates along with lower (2.5%) and upper (97.5%) percentiles are presented.

least 90% probability. However, the proportion of pixels which achieved elimination after 5 years reduced to 59% and 25% when the probability threshold was increased to 95% and 99%, respectively, illustrating that the policy is sensitive to uncertainty.
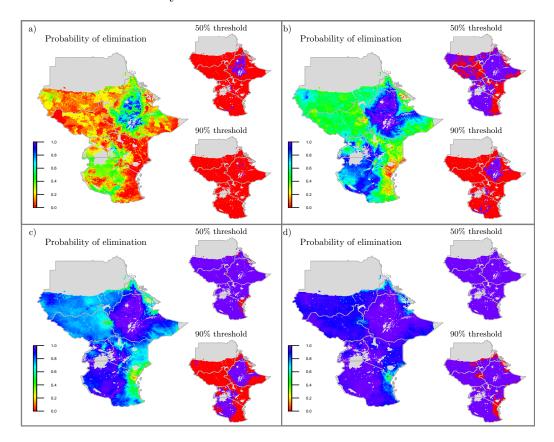

Figure 5: Probability of less than 1% prevalence after 5 years under: a) no intervention; annual MDA with coverage of b) 65%; c) 80% and d) biannual MDA at 65% coverage predicted at 5×5 km resolution, for Ethiopia, Sudan, South Sudan, Eritrea, Kenya, Tanzania and Uganda. Right of panels: Pixels that achieve elimination (blue) or do not achieve elimination (red) using different probability thresholds.

Finally, predictions of mf prevalence for the first and fifth year of intervention were summarized by calculating the estimated prevalence at each pixel, together with the 2.5$^{\text{th}}$ and 97.5$^{\text{th}}$ percentile in Figures D.2 and D.3 of the Appendix, for each of the four scenarios. Very similar observations were
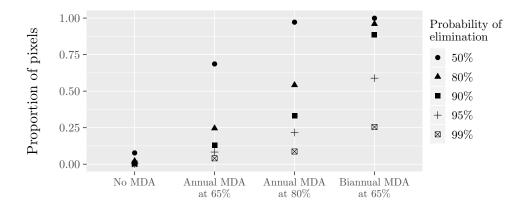
Figure 6: Proportion of pixels with prevalence less than 1% using different probability thresholds after 5 years under four intervention strategies.

made on the predictions of mf prevalence for the first 5 years of intervention.

5 Discussion

This study highlights the value of integrating geostatistical prevalence maps and transmission models for providing predictions on the impact of interventions aiming to eliminate transmission at a local scale. The main contribution is the development of new statistical tools through which existing research in mapping and predictive modelling are combined in a computationally efficient and flexible way which correctly accounts for uncertainty in these different techniques. Although we focus and apply our methodology on LF transmission, it can be applied to other infectious diseases.

We have shown that the current strategy of 5 annual rounds of MDA at 65% coverage will not be sufficient to eliminate disease in most areas. We also found that a change in the current MDA strategy, such as increasing the coverage and frequency of MDA, will be required if LF elimination is to be accelerated in East Africa. This suggests that it may be necessary to employ different enhanced intervention plans at a fine scale, according to the characteristics of each area, in order to achieve the WHO elimination targets.

However, for the results presented here we assumed that no interventions have been applied in East Africa prior to the prevalence survey. While this assumption is correct for most areas, MDA programs began to be implemented in a few districts of Africa since 2000 and in many more districts

thereafter. Therefore, one of the next steps will be to account for previous MDA programmes as spatio-temporal covariate information in the transmission model. Apart from MDA, insecticide treated bednets have been used in some countries (data can be extracted from the Malaria Atlas Project), the use of which has been shown to be an effective additional measure for control of the disease (Bockarie et al. 2009). Integrating geostatistical maps with transmission models with these additional covariates is more complicated as the simulations must include the appropriate historical interventions.

An additional challenge is the gap between reported and true coverage with an MDA. Where there are parasitological data against which to test the expected and achieved impact of reported coverages, they have been shown to be unreliable (Budge et al. 2016). This will pose a particular challenge to interpreting historic coverage and a challenge in communicating future projections. This work represents our initial framework and future research will be required to extend the methodology to capture these more complex settings.

A limitation of our statistical approach is that it doesn't capture the spatial correlations in the predictions, since each pixel is weighted independently to produce a marginal posterior for each pixel. This approach means that we lose the spatial autocorrelation that was captured in the original geostatistical model and, furthermore, that there is no way for nearby pixels to interact during the simulations, for example, to account for movements of humans or vectors. A more sophisticated approach would be to use the spatial autocorrelations from the geostatistical model, alongside any available movement or connectivity data, to define a transmission kernel that describes spatial spread. This kernel could be used within a single meta-population model describing the transmission dynamics across the whole map. At present, such an approach would be computational infeasible at the country scale, but may become possible in future through improvements in methodology and advances in parallelisation and cloud computing.

Acknowledgments

Appendix

In the appendix, we provide further details on the methodology and the transmission model as well as additional plots and results of the real data analysis.

A   Background information on the methodology

A.1   Proof of lemma on the change of measure

Proof. First we need to change the variables of the prior from $\boldsymbol{\theta}$ to the transformed variables $(p, \boldsymbol{q})$. This results in a transformed prior density $\pi'(p, \boldsymbol{q}) = \pi(\boldsymbol{\phi}^{-1}(p, \boldsymbol{q}))|J^{-1}|$, where J^{-1} is the determinant of the Jacobian matrix of partial derivatives of $\boldsymbol{\phi}^{-1}$. Next, we wish to apply the Radon-Nikodym derivative formula to change the measure over prevalences from $g(p)$ to $f(p)$, whilst keeping the measure over $\boldsymbol{q}$ unchanged. More formally, we rewrite $\pi'(p, \boldsymbol{q}) = \pi'(p)\pi'(p|\boldsymbol{q}) = g(p)\pi'(p|\boldsymbol{q})$ and change the whole measure from $g(p)\pi'(p|\boldsymbol{q})$ to $f(p)\pi'(p|\boldsymbol{q})$. This yields the new measure in terms of the transformed variables $h'(p, \boldsymbol{q}) = \frac{f(p)\pi'(p|\boldsymbol{q})}{g(p)\pi'(p|\boldsymbol{q})}\pi'(p, \boldsymbol{q}) = \frac{f(p)}{g(p)}\pi'(p, \boldsymbol{q})$. The final step is to transform back to the original variables $\boldsymbol{\theta}$.

$$h(\boldsymbol{\theta}) = h'(\phi(\boldsymbol{\theta}))|J|$$
$$= \frac{f(p(\boldsymbol{\theta}))}{g(p(\boldsymbol{\theta}))}\pi'(\phi(\boldsymbol{\theta}))|J|$$
$$= \frac{f(p(\boldsymbol{\theta}))}{g(p(\boldsymbol{\theta}))}\pi(\boldsymbol{\theta}).$$

□

A.2 Implementation considerations from the absolute continuity condition

The absolute continuity condition in Lemma 1, that f must be absolutely continuous with respect to g, means that whenever $g(p) = 0$ then we must also have $f(p) = 0$. In other words, when the prior probability of a prevalence is zero then the map measure of prevalence must also be zero. This has important considerations for implementing our method. For example, we find that our prior for the transmission model parameters is unlikely to produce any simulations with a prevalence of above 85%. However, prevalences this high do appear in the tails of the posterior distribution in the geostatistical model, especially when the amount of uncertainty is high. When the weighting is applied, there are no simulations to capture the high prevalence region, which leads to an underestimation of the mean prevalence, for example.

There are several potential approaches to ameliorate this. The simplest would be to mask out areas of the map where this occurs. However, there may be large areas of the map that have only a very small probability of producing a prevalence above 85% according to the geostatistical model. A second approach, which we could take if we believed that prevalences above 85% were a priori impossible, would be to remove any samples from the geostatistical map that went above 85%. This is effectively the same as applying the weighting naively, and leads to the mean and median of the prevalence distribution of the simulations being lower than the corresponding statistics from the geostatistical model. A final more sophisticated approach could be to shift the weight from high prevalences in the geostatistical map to the nearest prevalences available in the simulations. For example, by using the histogram-based empirical Radon-Nikodym derivative with a wide bin capturing all prevalences above 75%. Although this reduces the size of the underestimation, it substantially reduces the effective sample size, as a small number of simulations get very high weight.

A.3 Minimum discrepancy-based empirical Radon-Nikodym derivative

For each pixel, we would like to minimise the following distance between the empirical cumulative distribution functions (cdfs) of the posterior prevalences

and the weighted simulated prevalences:

$$\int_0^1 \left(F(x|\boldsymbol{d}) - H(x|\boldsymbol{w}^{(2)}) \right)^2 \, \mathrm{d}x,$$

with respect to the weights $\boldsymbol{w}^{(2)}$. First, the posterior and simulated prevalence samples are sorted in ascending order, such that $d_{(1)} \leq d_{(2)} \leq \ldots \leq d_{(M-1)} \leq d_{(M)}$ and $p_{(1)} \leq p_{(2)} \leq \ldots \leq p_{(J-1)} \leq p_{(J)}$. Similarly, we introduce the notation $w_{(j)}^{(2)}$ for the weights of simulation j which produces prevalence $p_{(j)}$, for $j = 1, 2, \ldots, J$. For clarity of notation, an example of two empirical cdfs is shown in the left panel of Figure A.1.
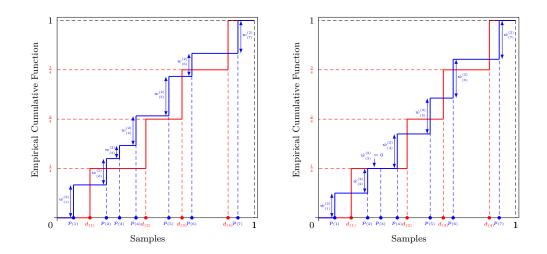

Figure A.1: An example of the empirical cumulative distribution functions of the posterior prevalences (red) and the weighted simulated prevalences (blue). The empirical cdf of the weighted simulated prevalence is shown in the left panel with arbitrary weights and in the right panel using the optimal weights that achieve the minimum distance between the two cdfs, as described in Section A.3.

24

Starting with $j = 1$, we have that:

$$\frac{\partial}{\partial w_{(1)}^{(2)}} \int_{p_{(1)}}^{p_{(2)}} \left(F(x|\boldsymbol{d}) - H(x|\boldsymbol{w}^{(2)}) \right)^2 \, \mathrm{d}x = \frac{\partial}{\partial w_{(1)}^{(2)}} \int_{p_{(1)}}^{p_{(2)}} \left(F(x|\boldsymbol{d}) - w_{(1)}^{(2)} \right)^2 \, \mathrm{d}x$$

$$= -2 \left[\int_{p_{(1)}}^{p_{(2)}} \left(F(x|\boldsymbol{d}) - w_{(1)}^{(2)} \right) \, \mathrm{d}x \right]$$

$$= -2 \left[\int_{p_{(1)}}^{p_{(2)}} F(x|\boldsymbol{d}) \, \mathrm{d}x - w_{(1)}^{(2)} \left(p_{(2)} - p_{(1)} \right) \right]. \tag{2}$$

At the minimum Equation (2) is equal to 0, therefore:

$$\hat{w}_{(1)}^{(2)} = \frac{\displaystyle\int_{p_{(1)}}^{p_{(2)}} F(x|\boldsymbol{d}) \, \mathrm{d}x}{p_{(2)} - p_{(1)}}.$$

Similarly, we differentiate the distance between the two cdfs with respect to $w_{(j)}^{(2)}$, for $j = 2, 3, \ldots, J - 1$ and by setting it equal to 0,

$$\frac{\partial}{\partial w_{(j)}^{(2)}} \int_{p_{(j)}}^{p_{(j+1)}} \left(F(x|\boldsymbol{d}) - \sum_{k=1}^{j-1} \hat{w}_{(k)}^{(2)} - w_{(j)}^{(2)} \right)^2 \, \mathrm{d}x$$

$$= -2 \left[\int_{p_{(j)}}^{p_{(j+1)}} F(x|\boldsymbol{d}) \, \mathrm{d}x - \left(p_{(j+1)} - p_{(j)} \right) \left(\sum_{k=1}^{j-1} \hat{w}_{(k)}^{(2)} + w_{(j)}^{(2)} \right) \right] = 0,$$

we obtain:

$$\hat{w}_{(j)}^{(2)} = \frac{\displaystyle\int_{p_{(j)}}^{p_{(j+1)}} F(x|\boldsymbol{d}) \, \mathrm{d}x}{p_{(J+1)} - p_{(J)}} - \sum_{k=1}^{j-1} \hat{w}_{(k)}^{(2)}.$$

Finally, to ensure that the weights sum to one, for $j = J$ we have that:

$$\hat{w}_{(J)}^{(2)} = 1 - \sum_{k=1}^{J-1} \hat{w}_{(k)}^{(2)}$$

As an example, in the right panel of Figure A.1 we provide the optimal weights $\hat{w}_{(j)}^{(2)}$, i.e. the ones with the minimum distance between the two empirical cdfs.

B   Toy example

Suppose that the prior distribution is $\pi(\theta_1, \theta_2) = 2$ if $0 < \theta_2 < \theta_1 < 1$ and zero otherwise. The prior support and marginal densities are shown in Figure B.1. For simplicity, assume that the transmission model has equilibrium prevalence given by $p(\theta_1, \theta_2) = \theta_1$ so that the induced prior over prevalences is the marginal for θ_1, ie. $g(p) = 2p$ for $0 < p < 1$. Further, suppose that we are given a pixel with prevalence measure $f(p) = 2(1 - p)$ for $0 < p < 1$. This challenging example allows us to assess how the methodology performs when there are few simulations with low weights in the area of high posterior probability close to $p = 0$.
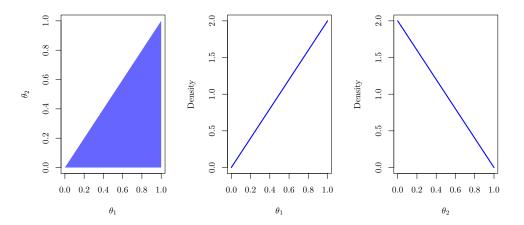

Figure B.1: Prior support and marginal densities for the parameters of the toy example in Section B.

By the Lemma 1 of the main text, the new measure over the parameter space is given by

$$
\begin{aligned}
h(\theta_1, \theta_2) &= \frac{f(p(\theta_1, \theta_2))}{g(p(\theta_1, \theta_2))} \pi(\theta_1, \theta_2) \\
&= \frac{f(\theta_1)}{g(\theta_1)} \times 2 \\
&= \frac{2(1 - \theta_1)}{2\theta_1} \times 2 \\
&= \frac{2(1 - \theta_1)}{\theta_1} \qquad \text{for } 0 < \theta_2 < \theta_1 < 1.
\end{aligned}
$$

26

First, we can verify that we get the correct marginal for $p = \theta_1$.

$$
\begin{aligned}
f(p) &= \int h(p, \theta_2) \, \mathrm{d}\theta_2 \\
&= \int_0^p 2 \frac{(1-p)}{p} \, \mathrm{d}\theta_2 \\
&= \left[\theta_2 \frac{2(1-p)}{p} \right]_{\theta_2 = p} \\
&= 2(1-p) \qquad \text{for } 0 < p < 1,
\end{aligned}
$$

as required.

Second, we can determine the new marginal density for θ_2.

$$
\begin{aligned}
h(\theta_2) &= \int h(\theta_1, \theta_2) \, \mathrm{d}\theta_1 \\
&= \int_{\theta_2}^1 \frac{2(1-\theta_1)}{\theta_1} \, \mathrm{d}\theta_1 \\
&= 2 \left[\log(\theta_1) - \theta_1 \right]_{\theta_1 = \theta_2}^1 \\
&= 2(\theta_2 - \log(\theta_2) - 1) \qquad \text{for } 0 < \theta_2 < 1. \qquad (3)
\end{aligned}
$$

B.1   Simulations

In this section, we perform a series of simulations to assess the accuracy and efficiency of the proposed method for recovering the distribution of pixel prevalence. Particular focus is given on how the method is affected as we vary the value of δ, the proposal distribution $q(\boldsymbol{\theta})$, and the empirical estimate of the Radon-Nikodym derivative.

B.1.1   Sensitivity analysis: Value of δ.

We first investigated the performance of the proposed method as a function of δ, where the observed pixel and simulated prevalence data are obtained from the toy model described in Section B. In particular, we generated $M = 2\,000$ pixel prevalence samples from Beta(1,2). We then generated $J = 2\,000$ samples from the joint prior distribution of parameters θ_1 and θ_2; first we drew $\theta_1 \sim \text{Beta}(2,1)$ and then $\theta_2 \mid \theta_1 \sim \theta_1 \times \text{Beta}(1,1)$. Therefore, the obtained simulated prevalence samples are draws from Beta(2,1), which is the marginal

prior for θ_1. We considered values for δ from 0.001 through to 0.1, increasing by 0.001 each time.

Accuracy was assessed by computing the Kolmogorov–Smirnov (KS) distance which is defined as the largest vertical distance between the two empirical cumulative distribution functions of the pixel prevalence and the weighted simulated prevalence. We also consider an additional measure of quantifying the distance between the two cdfs, termed as integrated squared distance and given by $\int_0^1 \left(F(x|\boldsymbol{d}) - H(x|\boldsymbol{w}^{(2)}) \right)^2 \, \mathrm{d}x$. The calculation of the two distances is repeated 100 times for each value of δ, using new pixel and simulated prevalence samples each time, in order to prevent biases occurring due to the simulating procedure. Results are shown in Figure B.2. We see that the performance of our method is affected by the value of δ. In the left panel of Figure B.2 we show the median KS distance along with the 95% credible interval, over the 100 realisations for each δ. Overall, the accuracy of the algorithm increases as δ grows from 0.001 to 0.033, where the median KS distance reaches its minimum value, and for δ higher than 0.033 a slight decrease is observed. The integrated squared distance, shown in the middle panel of Figure B.2, provides identical conclusions. Therefore, from now on we use the integrated squared difference to assess the accuracy of the method. An opposite pattern is observed in the efficiency of the method, as can be seen in the right panel of Figure B.2 where we show the effective sample size (ESS) as a function of δ. Note that despite having the highest ESS for $\delta = 0.001$, the method has the lowest accuracy as indicated by the highest distance between the two empirical distributions. This could be attributed to the fact that there are very few simulations in each estimate of density g (the denominator in Equation 1).

This motivated us to propose an appropriate value of δ depending on the simulated prevalence samples. More specifically, let $\delta_{k,j} = 2 \mid p_k - p_j \mid$, for $j, k = 1, 2, \ldots, J$. For each k, let $\delta_{k,(j)}$ to denote the sorted values with respect to j, such that $\delta_{k,(1)} \leq \delta_{k,(2)} \leq \ldots \leq \delta_{k,(J)}$. Note that $\delta_{k,(1)}$ is always equal to zero since for $k = j$, $\delta_{k,k} = 0$. Finally, the suggested value of δ, denoted by $\tilde{\delta}$, is given by $\tilde{\delta} = \max \left(\delta_{1,(3)}, \delta_{2,(3)}, \ldots, \delta_{J,(3)} \right)$, ensuring that at least 3 simulations were included in each estimate of g. In Figure B.3, we report the median $\tilde{\delta}$ (orange vertical line) along with the 95% credible interval, as obtained from the 100 simulations. Overall, the suggested value of δ achieves a good balance between efficiency and accuracy. In addition, the estimated weighted prevalence distribution of the suggested value of δ is provided in
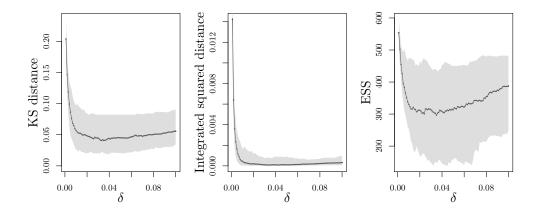
Figure B.2: Sensitivity to δ in the example where the parameters are drawn from their prior distribution. The left panel shows the Kolmogorov-Smirnov (KS), the middle panel the integrated squared distance and the right panel the ESS for $\delta = 0.001, 0.002, \ldots, 0.1$. Shaded areas correspond to the 95% credible interval.

the right panel of Figure 2(a) in the main text and shows that the method correctly reproduces the histogram of the true pixel prevalence (left panel), except maybe for prevalence values close to 0 where the histogram is a bit noisy. This is justified by the fact that there are only few simulations that have prevalence values between 0 and 0.02, as illustrated in the middle panel.

B.1.2 Sensitivity analysis: Proposal distribution.

In this section, we carry out a sensitivity analysis to assess the effect that the proposal distribution of the parameters has on the performance of the method. More specifically, instead of drawing θ_1 from its prior distribution Beta(2,1), we consider a uniform proposal $U(0, 1)$. This change results to a uniform distribution over the simulated prevalences. As before, we assess the performance of the method for different values of δ, shown in Figure B.3. Overall, we conclude that the performance of the method improves substantially when we move from the prior to the uniform proposal distribution over the prevalences. This is because there are no areas with strong posterior probability that have few simulations from the proposal. In particular, the median integrated squared distance between the two cdfs (left panel) ap-
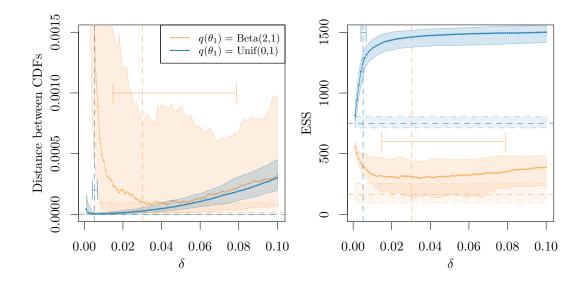
29

Figure B.3: Integrated squared distance between the two cumulative distribution functions (cdfs) (left panel) and effective sample size (ESS) (right panel) obtained from 100 simulated datasets, under different values of δ and choice of proposal distribution for parameter θ_1. Orange solid line represents a prevalence proposal distribution equal to marginal prior, i.e. Beta(2,1), whereas blue line corresponds to a $U(0,1)$ proposal. In both cases, the pixel prevalences were drawn from a Beta(1,2) distribution. Dashed horizontal lines correspond to the minimum possible distance (left panel) and its associated ESS (right panel). Shaded areas correspond to the 95% credible intervals for each estimate. Dashed vertical lines represent the suggested value of δ for each scenario of the proposal distribution considered, along with their upper and lower limits (error bars).

pears to be much lower in the latter compared to the former, and it is also associated with lower variability of the estimate. In addition, using a uniform proposal distribution leads to substantial increase in ESS (right panel), with the median ESS being at least 1.5 times higher for values of δ close to 0.001 and up to 3.5 times for larger values of δ. For reference, we also display the minimum discrepancy-based empirical Radon-Nikodym derivative, as described in Section A.3, and its associated ESS. Overall, we come close to reach the minimum possible integrated squared distance between the two cdfs, with larger ESS.

To further assess the accuracy of the method we compared the weighted posterior distribution of θ_2 with its true density, calculated in Equation (3). Results are shown in Figure B.4 for each choice of the proposal distribution. The simulation analysis illustrates that both proposals perform well in terms of recovering the target density of θ_2. Nevertheless, using a uniform proposal for the prevalence leads to a large improvement in the estimate.
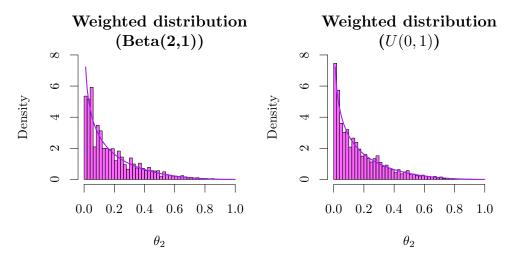

Figure B.4: The estimated weighted posterior distribution of θ_2 is compared to the target density (purple line), under different proposal distributions for the prevalence: Beta(2,1) (left panel) and $U(0,1)$ (right panel).

B.1.3 Sensitivity analysis: Empirical estimate of the Radon-Nikodym derivative.

In order to study the influence of the empirical Radon-Nikodym derivative (ERND) on the method performance, we repeated our simulation analysis for three different choices: (a) the proposed empirical estimate described in Section 2.2 referred as the distance-based derivative, (b) the histogram-based derivative in Section 2.5.1 and (c) the discrepancy-based derivative using the integrated squared distance (for more details see Section 2.5.2 of the manuscript). Results are given in Table B.1. Note that the proposed distance-based derivative is based on using the prevalences within $\tilde{\delta}/2$ of p_j and the histogram-based derivative is calculated by splitting the $[0, 1]$ interval into 100 equal bins.

As expected, the discrepancy-based ERND is optimal in terms of the distance between the two cdfs, but results in much lower ESS compared to the alternative derivatives. When there are few simulations of the proposal in the area of high posterior probability close to zero prevalence, the distance-based ERND performs better compared to the histogram-based ERND in terms of both integrated squared distance and ESS. When we move to a uniform simulated prevalence distribution, we find that histogram-based derivative outperforms the distance-based derivative, and is the one that scores highest in terms of ESS, followed in order by distance-based and discrepancy-based derivative. However, its performance depends on the choice of bins, since the relative weightings within each bin are unchanged. For example, Figure B.5 uses just 10 bins to illustrate that the distribution within each bin does not reflect the target distribution. Finally, we conclude that the performance of all the derivatives is greatly improved when we change the proposal from the prior to the uniform distribution over the prevalences.

Proposal distribution	ERND	Integrated squared distance ($\times 1000$)	ESS
Beta(2,1)	Distance	0.23734 (0.01742, 1.03220)	338 (222, 473)
	Histogram	0.46996 (0.02757, 1.56949)	335 (252, 453)
	Discrepancy	0.01647 (0.00335, 0.09844)	164 (96, 255)
U(0,1)	Distance	0.00408 (0.00157, 0.01799)	1248 (1161, 1330)
	Histogram	0.00214 (0.00175, 0.00292)	1347 (1253, 1429)
	Discrepancy	0.00025 (0.00021, 0.00029)	748 (713, 805)

Table B.1: Integrated squared distance between the two empirical cdfs (multiplied by 1000) and ESS, averaged over 100 replicates under three different empirical Radon-Nikodym derivatives (ERND) and different proposal distributions. The 95% credible intervals are shown in parentheses.
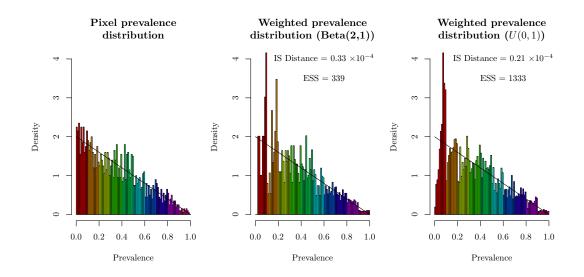
Figure B.5: The estimated weighted prevalence distributions obtained using the histogram-based empirical Radon-Nikodym derivative for 10 bins using a Beta(2,1) (middle panel) and $U(0,1)$ proposal (right panel) is compared to the true pixel prevalence distribution (left panel). Different colours correspond to prevalences in different bins. For reference we also provide the integrated squared (IS) distance and the ESS.

C   The mathematical model of LF transmission dynamics

C.1   TRANSFIL model description

We employed the mathematical model of lymphatic filariasis (LF) transmission TRANSFIL (Irvine et al. 2015) to carry out the analysis in this paper. The model is a stochastic microsimulation of individuals with worm burden, microfilariae (mf) and other demographic parameters in relation to age and exposure to risk. Humans are modelled individually with their own burden of male and female worms denoted by W_i^m and W_i^f, respectively. Mf concentration in the peripheral blood, denoted by M_i, is also modelled for each individual and according to the number of fertile female worms is increasing as well as decreasing at constant rate. The total mf density in the population contributes towards the current density of L3 larvae in the human-biting

mosquito population. Therefore, the model describes the dynamics of individual human, worm inside the host, mf inside the host and larvae inside the mosquito. A detailed description and mathematical formulation of the model are given in Irvine et al. (2015) and more recently in Smith et al. (2017), so here we provide a summary.

Worm dynamics

For each individual i, both male and female worms are added according to their bite risk b_i which is individually sampled from a gamma distribution with mean $= 1$ and shape parameter k. The rate at which an individual i acquires an adult worm, either female or male, is described by the following expression:

$$0.5\lambda b_i(V/H)\psi_1\psi_2 s_2 h(a),$$

where λ is the number of bites per mosquito, V/H is the ratio of vectors to hosts, ψ_1 is the probability that an L3 larvae leaves the host during a biting event, ψ_2 is the probability that the L3 enters the host, s_2 is the proportion of L3 within the host that develop into adult worms and $h(a)$ is the age-dependent biting rate which increases with body size to saturate at age nine. Finally, we assume that each worm has a constant death rate μ.

Mf dynamics

For an individual i, the dynamics of mf are given by the following expression:

$$\frac{\mathrm{d}M_i}{\mathrm{d}t} = \alpha W_i^f \mathbb{I}(W_i^m > 0) - \gamma M_i,$$

with α being the production rate of mf per worm, γ the constant death rate of mf and function $\mathbb{I}$ is one if there are male worms and zero if not. This expression follows from the fact that *W. bancrofti* is completely polygamous and therefore the mf production rate depends upon the number of female worms combined with the presence of at least one male in the human host.

Larvae dynamics

Larvae development occurs when mf entered the mosquito during a

blood meal form an infected host. There are two forms of this relationship depending on the genus of mosquito vectors as expressed below:

$$L(m) = \kappa_{s1} \left(1 - e^{-r_1 m/\kappa_{s1}}\right) \quad \text{for } Culex,$$

$$L(m) = \kappa_{s2} \left(1 - e^{-r_2 m/\kappa_{s2}}\right)^2 \quad \text{for } Anopheles,$$

where m is the concentration of mf per $20\mu L$ taken during a blood meal and r, κ_s are the parameters which are related to the functional form of the uptake curve (Gambhir and Michael 2008). The uptake of larvae is an average of mf concentration in the peripheral blood over all individuals weighted by their bite-risk b_i and is described by the following function:

$$\tilde{L} = \frac{\sum_i L(m_i) b_i}{\sum_i b_i},$$

giving the average number of larvae per mosquito. Using this expression we can calculate the averaged number of larvae taken up in the population $\tilde{L}$ as follows:

$$\frac{\mathrm{d}L}{\mathrm{d}t} = \lambda g \tilde{L} - (\sigma + \lambda \psi_1) L,$$

where λ is the number of bites per mosquito, g is the proportion of mosquitoes which pick up infection when biting an infected host, σ is the death rate of mosquitoes and ψ_1 is the proportion of L3 leaving the mosquito per bite.

Finally the equilibrium value for L3 in a mosquito is given by:

$$L^* = \frac{\lambda g \tilde{L}}{\sigma + \lambda \psi_1}.$$

Host dynamics

Each human begins with zero infection and as mentioned before, has a bite-rate of exposure drawn from a gamma distribution with mean $= 1$ and shape parameter k. The shape parameter defines how aggregated bites are amongst individuals and consequently defines the aggregation of infection amongst individuals (Irvine et al. 2015). Each individual has a rate of death τ that is assumed to be constant throughout an individual's lifetime with a cut-off at age 100.

A list of the basic model parameter values is provided in Table C.1, with specification of their source. In Table C.2 we provide the parameters that vary spatially across the study area, along with the distribution that are generated from.

Symbol	Definition	Value	Source
λ	Number of bites per mosquito	10 per month	Rajagopalan (1980), Subramanian et al. (1994)
$a_{\max}$	Age at which exposure to mosquitoes reaches its max level	20	Subramanian et al. (2004)
ψ_1	Proportion of L3 leaving mosquito per bite	0.414	Hairston and de Meillon (1968)
ψ_2	Proportion of L3 leaving mosquito that enter host	0.32	Ho and Ewert (1967)
s_2	Proportion of L3 entering host that develop into adult worms	0.00275	Norman et al. (2000), Stolk et al. (2008)
μ	Death rate of adult worms	0.0104 per month	Evans et al. (1993)
α	Production rate of Mf per worm	0.2 per month	Hairston and de Meillon (1968)
γ	Death rate of Mf	0.1 per month	Hairston and de Meillon (1968), Ottesen and Ramachandran (1995)
g	Proportion of mosquitoes which pick up infection when biting an infected host	0.37	Subramanian et al. (1998)
σ	Death rate of mosquitoes	5 per month	Ho and Ewert (1967)
$h(a)$	Parameter to adjust rate at which individuals of age a are bitten	Linear from 0 to 10, with max 1	Norman et al. (2000)
χ_1	Proportion of Mf killed for an individual MDA round using ALB and DEC	0.95	Ismail et al. (1998), Michael et al. (2004)

36

	Proportion of adult worm permanently sterilised during MDA round using ALB and DEC	0.55	Ismail et al. (1998), Michael et al. (2004)
κ_1			
ρ	Systematic adherence of MDA	0.35	Stolk et al. (2018b)

Table C.1: Description of the basic model parameters that are assumed to be fixed across the study area.
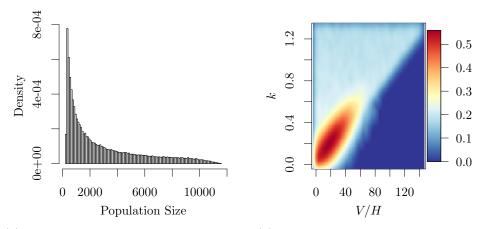
Symbol	Definition	Value
η	Size of population in simulation	Values are draw from the proposal distribution given in Figure C.1(a)
V/H	Ratio of number of vectors to hosts	Values are draw from the joint prior distribution given in Figure C.1(b)
k	Aggregation parameter of individual exposure to mosquitoes	Values are draw from the joint prior distribution given in Figure C.1(b)
α_{Imp}	Importation rate	Values are draw from a uniform prior distribution $\mathcal{U}(0, 0.0005)$
p_C	Coverage of MDA	65% or 80% depending on the future intervention assumptions

Table C.2: Description of the basic model parameters that are assumed to be varied across the study area.

C.2   Implementation details

To generate the require range of mf prevalences for the seven African countries that we considered, i.e Ethiopia, Sudan, South Sudan, Eritrea, Kenya, Tanzania and Uganda, we varied four parameters of the model; the population size, the vector to host ratio (V/H), the aggregation parameter of individual exposure to mosquitoes (k) and the importation rate (α_{Imp}). More specifically, for the importation rate we used a uniform prior distribution with minimum 0 and maximum 0.0005 (max $\frac{5}{10\,000}$ infections per month). We investigate different maximum values for the importation rate, and we chose the one which give us the desired results without driving the dynamics of the disease. In addition, the interventions reduce the prevalence over time, and therefore as years pass, we decrease the importation rate after intervention in proportion to the reduction in prevalence seen in pilot simulations. More specifically, we produce 2 000 simulations with constant importation rate over five years of MDA. Then, for our main set of simulations, we adjust the importation rate according to how the prevalence changed in our pilot runs after the intervention was applied.

Parameters V/H and k were drawn from a range of plausible values based on previously analysed data (Irvine et al. 2015, 2017, Smith et al. 2017). The graphical representation of their prior distribution is shown in the left panel of Figure C.1. We then generate 100 000 parameter vectors by randomly sampled from these prior distributions of parameters V/H, k and α_{Imp} and from the proposal distribution of the population size as described in Section 4.3 and showed in the right panel of Figure C.1. These samples were then used to generate 100 000 simulations from the model for each scenario specified in the main text, using the exact years and coverage of the MDA treatments. Our simulations here are focused on areas with anopheles as the dominant vector species. Finally, we set $\delta = 0.01$.
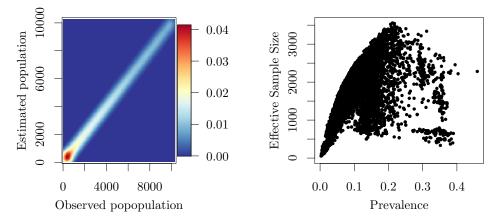
(a) Proposal density of the population sizes.

(b) Joint prior distribution on the vector to host ratio (V/H) and aggregation parameter of individual exposure to mosquitoes (k).

Figure C.1: Distributions assigned to parameters that vary spatially across the study area.

D   Additional results

In this section we provide additional results for the analysis of the LF data in East Africa of Section 4.3 of the main manuscript. The performance of the method described in Section 2 was assessed by comparing the observed and the estimated number of people at each pixel in the left panel of Figure D.1. In the right panel of Figure D.1, we further evaluated the performance of our method by plotting effective sample size per pixel. Figures D.2 and D.3 illustrate future predictions of the prevalence under different control scenarios for the first and fifth year.

(a) Comparison of the estimated number of people per pixel with the observed value.

(b) Effective sample size against the estimated median prevalence for each pixel.

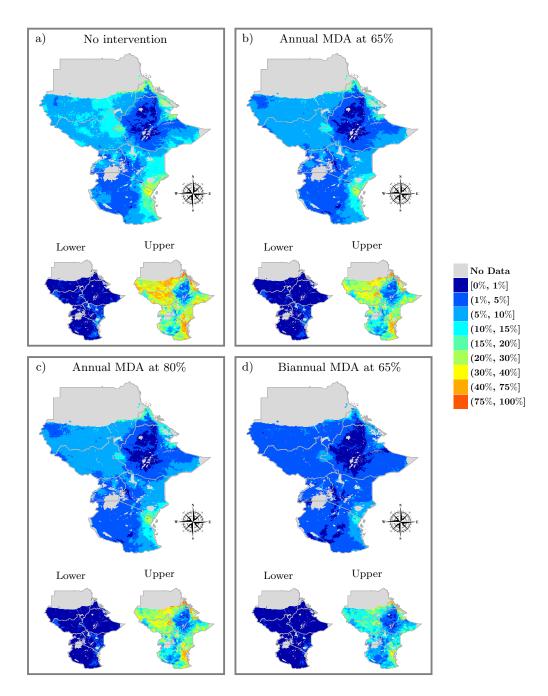Figure D.1: Performance assessment of our method.

Figure D.2: Prevalence after 1 year under: a) no intervention; annual MDA with coverage of b) 65%; c) 80% and d) biannual MDA at 65% coverage predicted at 5×5 km resolution. Point estimates along with lower (2.5%) and upper (97.5%) percentiles are presented.
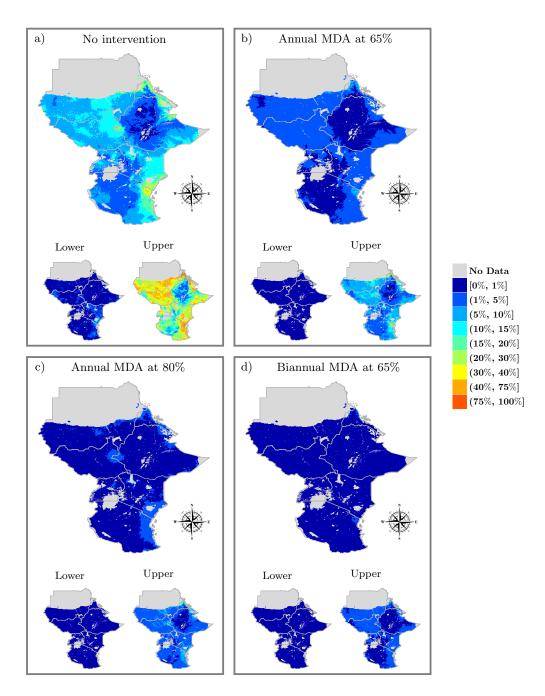
41

Figure D.3: Prevalence after 5 years under: a) no intervention; annual MDA with coverage of b) 65%; c) 80% and d) biannual MDA at 65% coverage predicted at 5×5 km resolution. Point estimates along with lower (2.5%) and upper (97.5%) percentiles are presented.

References

Alley, E. S., Plaisier, A. P., Boatin, B. A., Dadzie, K. Y., Remme, J., Zerbo, G., Samba, E. M., 1994. The impact of five years of annual ivermectin treatment on skin microfilarial loads in the onchocerciasis focus of Asubende, Ghana. Transactions of the Royal Society of Tropical Medicine and Hygiene 88 (5), 581–584.

Beaumont, M. A., Zhang, W., Balding, D. J., 2002. Approximate Bayesian computation in population genetics. Genetics 162 (4), 2025–2035.

Bhatt, S., Weiss, D. J., Cameron, E., Bisanzio, D., Mappin, B., Dalrymple, U., Battle, K. E., Moyes, C. L., Henry, A., Eckhoff, P. A., et al., 2015. The effect of malaria control on Plasmodium falciparum in Africa between 2000 and 2015. Nature 526, 207–211.

Billingsley, P., 1995. Measure and Probability.

Bockarie, M. J., Pedersen, E. M., White, G. B., Michael, E., 2009. Role of vector control in the global program to eliminate lymphatic filariasis. Annual Review of Entomology 54, 469–487.

Budge, P. J., Sognikin, E., Akosa, A., Mathieu, E. M., Deming, M., 2016. Accuracy of coverage survey recall following an integrated mass drug administration for lymphatic filariasis, schistosomiasis, and soil-transmitted helminthiasis. PLoS Neglected Tropical Diseases 10 (1), e0004358.

Cornuet, J.-M., Marin, J.-M., Mira, A., Robert, C. P., 2012. Adaptive multiple importance sampling. Scandinavian Journal of Statistics 39 (4), 798–812.

Deardon, R., Brooks, S. P., Grenfell, B. T., Keeling, M. J., Tildesley, M. J., Savill, N. J., Shaw, D. J., Woolhouse, M. E., 2010. Inference for individual-level models of infectious diseases in large populations. Statistica Sinica 20 (1), 239.

Diggle, P. J., Ribeiro, P. J., 2007. Model-based Geostatistics (Springer Series in Statistics).

Evans, D. B., Gelband, H., Vlassoff, C., 1993. Social and economic factors and the control of lymphatic filariasis: a review. Acta tropica 53 (1), 1–26.

Ferguson, N. M., Cummings, D. A., Cauchemez, S., Fraser, C., Riley, S., Meeyai, A., Iamsirithaworn, S., Burke, D. S., 2005. Strategies for containing an emerging influenza pandemic in Southeast Asia. Nature 437 (7056), 209–214.

Gambhir, M., Michael, E., 2008. Complex ecological dynamics and eradicability of the vector borne macroparasitic disease, lymphatic filariasis. PLoS One 3 (8), e2874.

Gibson, G. J., 1997. Markov chain Monte Carlo methods for fitting spatiotemporal stochastic models in plant epidemiology. Journal of the Royal Statistical Society: Series C (Applied Statistics) 46 (2), 215–233.

Giorgi, E., Diggle, P., Snow, R., Noor, A., 2018. Geostatistical methods for disease mapping and visualisation using data from spatio-temporally referenced prevalence surveys. International Statistical Review 86 (3), 571–597.

Goldie, C. M., Maller, R. A., 1999. Generalized densities of order statistics. Statistica Neerlandica 53 (2), 222–246.

Griffin, J. T., Hollingsworth, T. D., Okell, L. C., Churcher, T. S., White, M., Hinsley, W., Bousema, T., Drakeley, C. J., Ferguson, N. M., Basáñez, M.-G., et al., 2010. Reducing Plasmodium falciparum malaria transmission in Africa: a model-based evaluation of intervention strategies. PLoS medicine 7 (8), e1000324.

Hairston, N. G., de Meillon, B., 1968. On the inefficiency of transmission of Wuchereria bancrofti from mosquito to human host. Bulletin of the World Health Organization 38 (6), 935.

Hay, S. I., Guerra, C. A., Gething, P. W., Patil, A. P., Tatem, A. J., Noor, A. M., Kabaria, C. W., Manh, B. H., Elyazar, I. R. F., Brooker, S., et al., 2009. A world malaria map: Plasmodium falciparum endemicity in 2007. PLoS medicine 6 (3), e1000048.

Heesterbeek, H., Anderson, R. M., Andreasen, V., Bansal, S., De Angelis, D., Dye, C., Eames, K. T. D., Edmunds, W. J., Frost, S. D. W., Funk, S., et al., 2015. Modeling infectious disease dynamics in the complex landscape of global health. Science 347 (6227), aaa4339.

Ho, B. C., Ewert, A., 1967. Experimental transmission of filarial larvae in relation to feeding behaviour of the mosquito vectors. Transactions of the Royal Society of Tropical Medicine and Hygiene 61 (5), 663–666.

Hollingsworth, T. D., 2018. Counting down the 2020 goals for 9 neglected tropical diseases: What have we learned from quantitative analysis and transmission modeling? Clinical Infectious Diseases 66 (suppl_4), S237–S244.

Irvine, M. A., Reimer, L. J., Njenga, S. M., Gunawardena, S., Kelly-Hope, L., Bockarie, M., Hollingsworth, T. D., 2015. Modelling strategies to break transmission of lymphatic filariasis - aggregation, adherence and vector competence greatly alter elimination. Parasites and Vectors 8, 547.

Irvine, M. A., Stolk, W. A., Smith, M. E., Subramanian, S., Singh, B. K., Weil, G. J., Michael, E., Hollingsworth, T. D., 2017. Effectiveness of a triple-drug regimen for global elimination of lymphatic filariasis: a modelling study. The Lancet Infectious Diseases 17 (4), 451–458.

Ismail, M. M., Jayakody, R. L., Weil, G. J., Nirmalan, N., Jayasinghe, K. S. A., Abeyewickrema, W., Sheriff, M. H. R., Rajaratnam, H. N., Amarasekera, N. D. D. M., De Silva, D. C. L., et al., 1998. Efficacy of single dose combinations of albendazole, ivermectin and diethylcarbamazine for the treatment of bancroftian filariasis. Transactions of the Royal Society of Tropical Medicine and Hygiene 92 (1), 94–97.

Keeling, M. J., Woolhouse, M. E., Shaw, D. J., Matthews, L., Chase-Topping, M., Haydon, D. T., Cornell, S. J., Kappey, J., Wilesmith, J., Grenfell, B. T., 2001. Dynamics of the 2001 UK foot and mouth epidemic: stochastic dispersal in a heterogeneous landscape. Science 294 (5543), 813–817.

Michael, E., Malecela-Lazaro, M. N., Simonsen, P. E., Pedersen, E. M., Barker, G., Kumar, A., Kazura, J. W., 2004. Mathematical modelling and the control of lymphatic filariasis. The Lancet Infectious Diseases 4 (4), 223–234.

Moraga, P., Cano, J., Baggaley, R. F., Gyapong, J. O., Njenga, S. M., Nikolay, B., Davies, E., Rebollo, M. P., Pullan, R. L., Bockarie, M. J., et al., 2015. Modelling the distribution and transmission intensity of lymphatic filariasis in sub-Saharan Africa prior to scaling up interventions: integrated

use of geostatistical and mathematical modelling. Parasites and Vectors 8, 560.

Norman, R. A., Chan, M.-S., Srividya, A., Pani, S. P., Ramaiah, K. D., Vanamail, P., Michael, E., Das, P. K., Bundy, D. A. P., 2000. EPIFIL: the development of an age-structured model for describing the transmission dynamics and control of lymphatic filariasis. Epidemiology and Infection 124 (3), 529–541.

O'Hanlon, S. J., Slater, H. C., Cheke, R. A., Boatin, B. A., Coffeng, L. E., Pion, S. D. S., Boussinesq, M., Zouré, H. G. M., Stolk, W. A., Basáñez, M.-G., 2016. Model-based geostatistical mapping of the prevalence of Onchocerca volvulus in West Africa. PLoS Neglected Tropical Diseases 10 (1), e0004328.

Ottesen, E. A., Ramachandran, C. P., 1995. Lymphatic filariasis infection and disease: Control strategies. Parasitology Today 11 (4), 129–131.

Plaisier, A. P., Van Oortmarssen, G. J., Remme, J., Alley, E. S., Habbema, J. D., 1991. The risk and dynamics of onchocerciasis recrudescence after cessation of vector control. Bulletin of the World Health Organization 69 (2), 169–178.

Pullan, R. L., Sturrock, H. J. W., Magalhaes, R. J. S., Clements, A. C. A., Brooker, S. J., 2012. Spatial parasite ecology and epidemiology: a review of methods and applications. Parasitology 139 (14), 1870–1887.

Rajagopalan, P. K., 1980. Population dynamics of Culex pipiens fatigans, the filariasis vector, in Pondicherry: Influence of climate and environment. Proceedings of the Indian National Science Academy 6, 745–752.

Ramaiah, K. D., Ottesen, E. A., 2014. Progress and impact of 13 years of the global programme to eliminate lymphatic filariasis on reducing the burden of filarial disease. PLoS Neglected Tropical Diseases 8 (11), e3319.

Slater, H., Michael, E., 2013. Mapping, Bayesian geostatistical analysis and spatial prediction of lymphatic filariasis prevalence in Africa. PloS one 8 (8), e71574.

Smith, M. E., Singh, B. K., Irvine, M. A., Stolk, W. A., Subramanian, S., Hollingsworth, T. D., Michael, E., 2017. Predicting lymphatic filariasis transmission and elimination dynamics using a multi-model ensemble framework. Epidemics 18, 16–28.

Stensgaard, A.-S., Vounatsou, P., Onapa, A. W., Simonsen, P. E., Pedersen, E. M., Rahbek, C., Kristensen, T. K., 2011. Bayesian geostatistical modelling of malaria and lymphatic filariasis infections in Uganda: predictors of risk and geographical patterns of co-endemicity. Malaria Journal 10, 298.

Stolk, W. A., De Vlas, S. J., Borsboom, G. J. J. M., Habbema, J. D. F., 2008. LYMFASIM, a simulation model for predicting the impact of lymphatic filariasis control: quantification for African villages. Parasitology 135 (13), 1583–1598.

Stolk, W. A., Prada, J. M., Smith, M. E., Kontoroupis, P., De Vos, A. S., Touloupou, P., Irvine, M. A., Brown, P., Subramanian, S., Kloek, M., et al., 2018a. Are alternative strategies required to accelerate the global elimination of lymphatic filariasis? insights from mathematical models. Clinical infectious diseases 66 (Supplement_4), S260–S266.

Stolk, W. A., Prada, J. M., Smith, M. E., Kontoroupis, P., De Vos, A. S., Touloupou, P., Irvine, M. A., Brown, P., Subramanian, S., Kloek, M., et al., 2018b. Are alternative strategies required to accelerate the global elimination of lymphatic filariasis? Insights from mathematical models. Clinical Infectious Diseases 66, S260–S266.

Sturrock, H. J. W., Gething, P. W., Clements, A. C. A., Brooker, S., 2010. Optimal survey designs for targeting chemotherapy against soil-transmitted helminths: effect of spatial heterogeneity and cost-efficiency of sampling. The American Journal of Tropical Medicine and Hygiene 82 (6), 1079–1087.

Subramanian, S., Krishnamoorthy, K., Ramaiah, K. D., Habbema, J. D. F., Das, P. K., Plaisier, A. P., 1998. The relationship between microfilarial load in the human host and uptake and development of Wuchereria bancrofti microfilariae by Culex quinquefasciatus: a study under natural conditions. Parasitology 116 (3), 243–255.

Subramanian, S., Manoharan, A., Ramaiah, K. D., Das, P. K., 1994. Rates of acquisition and loss of Wuchereria bancrofti infection in Culex quinquefasciatus. The American Journal of Tropical Medicine and Hygiene 51 (2), 244–249.

Subramanian, S., Stolk, W. A., Ramaiah, K. D., Plaisier, A. P., Krishnamoorthy, K., Van Oortmarssen, G. J., Amalraj, D. D., Habbema, J. D. F., Das, P. K., 2004. The dynamics of Wuchereria bancrofti infection: a model-based analysis of longitudinal data from Pondicherry, India. Parasitology 128 (5), 467–482.

Tatem, A. J., Smith, D. L., Gething, P. W., Kabaria, C. W., Snow, R. W., Hay, S. I., 2010. Ranking of elimination feasibility between malaria-endemic countries. The Lancet 376 (9752), 1579–1591.

Tekle, A. H., Zouré, H. G. M., Noma, M., Boussinesq, M., Coffeng, L. E., Stolk, W. A., Remme, J. H. F., 2016. Progress towards onchocerciasis elimination in the participating countries of the African Programme for Onchocerciasis Control: epidemiological evaluation results. Infectious Diseases of Poverty 5 (1), 66.

Tildesley, M. J., Bessell, P. R., Keeling, M. J., Woolhouse, M. E., 2009. The role of pre-emptive culling in the control of foot-and-mouth disease. Proceedings of the Royal Society B: Biological Sciences 276 (1671), 3239–3248.

Wakefield, J., Lyons, H., 2010. Spatial aggregation and the ecological fallacy. Handbook of Spatial Statistics, 541–558.

WHO, 2012. Accelerating work to overcome the global impact of neglected tropical diseases: A roadmap for implementation Geneva (WHO/HTM/NTD/2012).

Wilkinson, R. D., 2013. Approximate bayesian computation (abc) gives exact results under the assumption of model error. Statistical applications in genetics and molecular biology 12 (2), 129–141.

Worldpop, 2010. The AfriPop demography project (www.worldpop.org.uk).