Proceedings of the Twenty-Ninth AAAI Conference on Innovative Applications (IAAI-17))

Crowdsensing Air Quality with Camera-Enabled Mobile Devices

Zhengxiang Pan,[1] Han Yu,[1] Chunyan Miao,[1,2] Cyril Leung[3,1]
[1]Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY),
Nanyang Technological University (NTU), Singapore
[2]School of Computer Science and Engineering (SCSE), NTU, Singapore
[3]Department of Electrical and Computer Engineering, University of British Columbia (UBC), Vancouver, BC, Canada
{panz0012, han.yu, ascymiao}@ntu.edu.sg, cleung@ece.ubc.ca

Abstract

Crowdsensing of air quality is a useful way to improve public awareness and supplement local air quality monitoring data. However, current air quality monitoring approaches are either too sophisticated, costly or bulky to be used effectively by the mass. In this paper, we describe AirTick, a mobile app that can turn any camera enabled smart mobile device into an air quality sensor, thereby enabling crowdsensing of air quality. AirTick leverages image analytics and deep learning techniques to produce accurate estimates of air quality following the Pollutant Standards Index (PSI). We report the results of an initial experimental and empirical evaluations of AirTick. The AirTick tool has been shown to achieve, on average, 87% accuracy in day time operation and 75% accuracy in night time operation. Feedbacks from 100 test users indicate that they perceive AirTick to be highly useful and easy to use. Our results provide a strong positive case for the benefits of applying artificial intelligence techniques for convenient and scalable crowdsensing of air quality.

Introduction

One of the most serious environmental challenges facing today's world is air pollution. Maintaining good air quality is important for public health. Every year, over 7 million people die as a result of air pollution (WHO 2014). This number is even higher than the number of premature deaths linked to dirty water or poor sanitation. Pollution by airborne particles as a result of persistent wild fires or industry pollution, often referred to as "haze" in layman terms, has been affecting major population centres in Southeast Asia, India and China for years. In Indonesia alone, more than 28 million people are affected by haze in 2015, with an estimated economic cost of US$47 billion (Chan 2015). A record high Pollutant Standards Index (PSI) (NEA 2014) value of 3,300 occurred in Oct 2015 in Kalimantan (Soeriaatmadja 2015). The occurrence of haze has adversely affected public health including elevating the risks of respiratory and cardiovascular diseases (Anderson, Thundiyil, and Stolbach 2012). In order to combat haze, countries are looking into motivating population participation in air quality monitoring.

Crowdsourcing is an emerging Internet-empowered phenomenon where the time, effort and resources from many individuals are mobilized towards specific goals (Silberzahn and Uhlmann 2015; Yu et al. 2016). It is a promising approach to help mobilize the collective effort of a society to monitor, and possibly improve, air quality over large geographic regions. In order for this to happen, a suitable air quality monitoring tool is needed. Currently, popular approaches of air quality monitoring include tapered element oscillating microbalance (Ruppecht, Meyer, and Patashnick 1992), black smoke measurement (Muir and Laxen 1995), and filter-based gravimetric methods (Hauck et al. 2004). However, these methods tend to be either too costly, too sophisticated, or too bulky for most people to use on the go. A simple, portable, scalable, and low cost way to monitor air pollution is needed to support crowdsensing of air quality, and to supplement local air quality data coverage.

Air pollution is often characterized by reduced visibility due to the scattering of light by airborne particles. Our naked eyes can roughly distinguish how hazy the sky is, but are generally unable to quantify the degree of air pollution. In this paper, we present *AirTick* – a mobile app which operates based on this principle with capability to quantify air pollution (following the Pollutants Standard Index (PSI) system adopted by the National Environment Agency (NEA) of Singapore) by leveraging on image analytics and deep learning technologies (Figure 1). AirTick is equipped the Adaptive Transmission Map (ATM) image analytics approach which is based on (Levin, Lischinski, and Weiss 2008). When a user wants to know the air quality around his immediate vicinity, he can simply turn on the AirTick app and take a photo of the outdoor environment surrounding him. ATM then analyzes the photo, extracts the haziness component from the image, and passes this component to the Deep Neural network Air quality estimator (DNA) – to produce an estimated PSI value for the user. DNA is powered by a Deep Boltzmann Machine (DBM) (Hinton and Salakhutdinov 2006). By leveraging people's familiarity (Pan et al. 2015) with using camera phones, AirTick can turn smart mobile devices into an air quality sensor, thereby providing a simple, portable, scalable, and low cost tool to convert a ordinary users into a network of sensing agents (Yu, Shen, and Leung 2013).

We have constructed a large-scale dataset consisting of

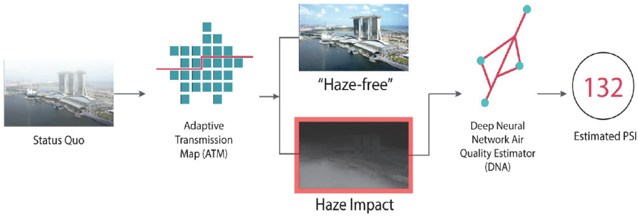

Figure 1: The AirTick technical framework.

1,000 photos taken in Singapore from 2014 to 2016 with the corresponding PSI values at the time obtained from the NEA historical PSI value online archive. This dataset is used to train and evaluate the DNA approach. On average, AirTick can achieve close to 87% accuracy for photos taken during day time, and close to 75% accuracy for photos taken at night. A user study involving 100 Singapore residents shows that they found AirTick to be highly useful and easy to use. Our results provide a strong case for the benefits of applying artificial intelligence (AI) techniques for convenient and scalable crowdsensing of air quality. Although we do not expect that AirTick could completely replace gold standard air quality monitoring techniques, the fact that it is mobile device ready and easy to use can help raise public awareness of air quality and provide timely recommendations for people with respiratory and other related conditions to take shelter from air pollution when necessary.

Related Work

Crowdsensing for air quality monitoring is starting to emerge in recent years due to increased public awareness and the availability of smaller and more affordable sensors. In (Jutzeler, Li, and Faltings 2014), the authors proposed a novel region-based Gaussian process model for estimating urban air pollution dispersion. The approach focuses on data aggregation and extraction of patterns from the air quality data contributed by the crowd. The approach partitions an urban area into regions which are assumed to be experiencing homogeneous pollutant emissions. By taking into account the spatial and land-use characteristics of the regions, the average air pollution level within a region is estimated. The new model is useful for applications such as exposure assessment and anomaly detection. However, it does not address the lack of a suitable air quality monitoring tool to enable crowdsensing, which is the focus of this paper.

(Liu et al. 2016) took a similar approach to ours in terms of estimating air quality based on analyzing outdoor photos. The proposed method utilizes six image features together with additional information such as the position of the sun, date, time, geographic information and weather conditions, etc., to estimated the amount of $PM_{2.5}$ particles in the air. Experimental results have shown that the image analysis method is able to estimate the $PM_{2.5}$ index accurately. Nevertheless, the method relies on the building regions of interest to be manually labelled in order to operate effectively. This step requires the users to precisely label the buildings in the photos they have taken, which incurs significant overhead. Furthermore, the additional information required by the method on top of the photos and labels of buildings may not always be available, especially in outdoor locations without Internet access.

In comparison, AirTick does not require manual input from the user or additional information in order to produce air quality estimates with good accuracy. This makes AirTick more suitable for large-scale usage as an easy to use personal air quality sensor.

The AirTick Prototype System

In this section, we present the design and implementation details regarding AirTick – the proposed mobile app which turns anyone's mobile device (e.g., hand phone or Pad) into an air quality sensor. AirTick consists of two main technological components, each addressing one of the following two major research challenges:

1. How to extract the haziness from a single image without knowing what the corresponding "haze-free" version of the image looks like?
2. How to convert the haziness extracted from an image into a PSI value accurately?

Extracting Haziness from a Photo

As shown in Figure 2, when taking a photo outdoors, object features gradually appear lighter and fade as they get closer to the horizon. Only a proportion of the reflected light reaches the camera as a result of absorption by the atmosphere. In addition, the light is mixed with the airlight colour vector, and due to its scattering effects, the scene colour is shifted. In general, a hazy image can be modelled as:

$$I(w,h) = e^{-\epsilon d(w,h)}R(w,h) + (1 - e^{-\epsilon d(w,h)})L_{\infty} \quad (1)$$

where $I(w,h)$ is the observed hazy image (w and h denotes the width and height of the image in pixels, respectively); ϵ is the atmospheric attenuation coefficient; L_{∞} is the atmospheric light; $d(w,h)$ is the scene depth; and $R(w,h)$ is the scene radiance (a.k.a. the haze-free image). Atmospheric attenuation exponentially degrades the scene radiance according to $d(w,h)$. Atmospheric light is a white atmospheric veil interfering with the radiance and reducing the visibility.

The term $e^{-\epsilon d(w,h)}$ is referred to as the transmission. Based on Eq. (1), $e^{-\epsilon d(w,h)}$ can be expressed as:

$$e^{-\epsilon d(w,h)} = \alpha I(w,h) + \beta \quad (2)$$

where $\alpha = \frac{1}{R(w,h)-L_{\infty}}$ and $\beta = \frac{L_{\infty}}{R(w,h)-L_{\infty}}$.

As L_{∞} can be determined using empirical techniques, in order to estimate the haziness of a given image, the two unknowns $d(w,h)$ and $R(w,h)$ need to be solved simultaneously. With only one photo of the environment, there is no ground truth image about how the place looks like without haze. Thus, the ATM approach divides an image into many small local windows and leverages the assumption that the scene radiance is approximately constant within a local window. By finding the image transmission map and radiance values which can reduce the haziness in the image to a predetermined level, ATM quantifies the haziness in a given image. Based on the principles from (Levin, Lischinski, and

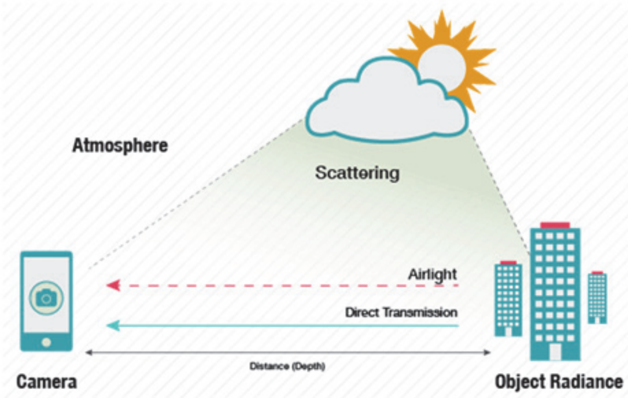

Figure 2: The optical model of a hazy outdoor photo.

Weiss 2008), $d(w,h)$ and $R(w,h)$ can be found by minimizing the following objective function.

Minimize:

$$\sum_{i \in I(w,h)} \left[\sum_{j \in \omega_i} \left(e_j^{-\epsilon d(w,h)} - (\alpha_i I(w,h)_j + \beta_i) \right)^2 + \psi \alpha_i^2 \right] \tag{3}$$

Subject to:

$$e_j^{-\epsilon d(w,h)} \geqslant 1 - \frac{I(w,h)_j}{L_\infty}, \forall j. \tag{4}$$

ω_i is a local window around a pixel i in the image. ψ is a regularization parameter for α_i. Constraint (4) is there to ensure that the resulting $R(w,h)$ image obtained does not contain negative pixel values.

To minimize Eq. (3), ATM moves a small $n \times n$ pixel local window through each pixel of the image to solve the optimization and produce pixel values in $d(w,h)$ and $R(w,h)$. The solution for the objective function is applied to each colour channel of a given image separately to extract the haziness from the image. In order for the local window to be small enough such that the assumption that the scene radiance is approximately constant is upheld, we set $n = 3$ in the AirTick prototype system.

Converting the Extracted Haziness into PSI Values

With haziness extracted from a given image by ATM, AirTick passes the haziness information to DNA to learn to associate given haziness matrices with PSI values. DNA is designed based on the Boltzmann Machine (BM). The original BM was proposed in (Hinton and Salakhutdinov 2006). It is a neural network of symmetrically coupled stochastic binary nodes. It contains a set of D visible nodes $\mathbf{v} \in \{0,1\}^D$, and a set of P hidden nodes $\mathbf{h} \in \{0,1\}^P$. The effectiveness (a.k.a. energy function) of a BM with a given state $\{\mathbf{v}, \mathbf{h}\}$ is defined as:

$$E(\mathbf{v}, \mathbf{h}; \theta) \triangleq -\frac{1}{2}\mathbf{v}^T \mathbf{L} \mathbf{v} - \frac{1}{2}\mathbf{h}^T \mathbf{J} \mathbf{h} - \frac{1}{2}\mathbf{v}^T \mathbf{W} \mathbf{h} \tag{5}$$

where $\theta = \{\mathbf{L}, \mathbf{J}, \mathbf{W}\}$ are the model parameters which need to be trained. $\mathbf{L}$, $\mathbf{J}$ and $\mathbf{W}$ correspond to the visible-to-visible, hidden-to-hidden and visible-to-hidden symmetric interaction terms, respectively.

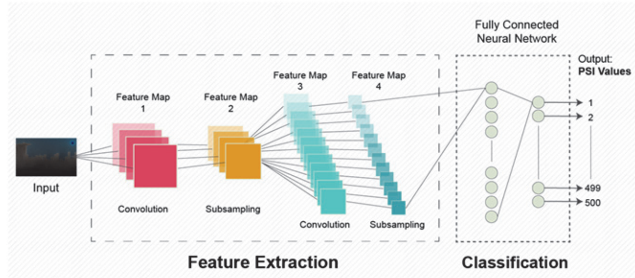

Figure 3: The DNA technical framework.

DNA learns in such a way that the hidden activities of one level of the structure are used as the data for training the higher level. The number of visible nodes involved in the model will determine the granularity of the estimated air quality levels. The larger the number of visible nodes, the more fine-grained the estimations can be. However, increasing the number of visible nodes will also incur higher cost for training the model. Due to high complexity typical of a fully connected BM, we design an alternative approach to avoid directly learning the large number of parameters involved. Instead, we focus on learning a deep multi-layer BM as shown in Figure 3, in which each layer captures complex and higher-order correlations among the activities of hidden nodes in the layer below.

When a user uses AirTick (i.e., taking a photo of the environment), the time, location and camera orientation information is automatically logged by AirTick. When the user's mobile device is connected to the Internet, the photo, together with these pieces of information, are uploaded to the AirTick server. The server then searches for the actual PSI value from NEA's historical PSI value online archive at the time the photo was taken, and incorporates the new photo and meta-data into our training dataset to improve DNA.

Psychological Priming

In addition to producing useful air quality estimations based on images, we also intend to use AirTick to promote mass awareness and citizen stewardship for maintaining good air quality. Thus, in the user interface design of AirTick, we infused the component of psychological priming. Priming refers to an implicit memory effect in which exposure to one perceptual stimulus influences people's responses to another stimulus (Meyer and Schvaneveldt 1971). As the ATM module of AirTick produces both the haziness component and an image with the haziness component removed based on a given image (as shown in Figures 4 and 5), we infuse these two types of output as part of the AirTick user interface to achieve psychological priming.

Figure 6 shows the user interface of AirTick. Once a user takes a photo of the environment using AirTick, the haziness extracted from the image and the original image with haze removed are displayed next to the original image. In this way, the app provides the user with a visual feedback on

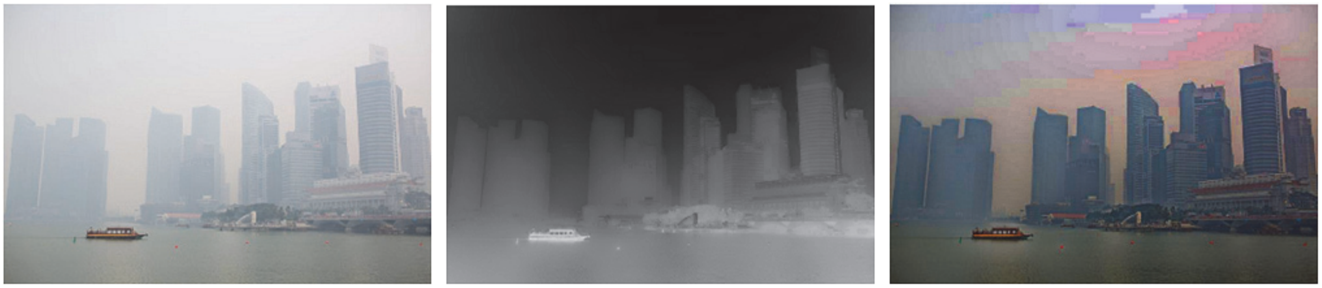

(a) The original image (b) The haze component extracted by ATM (c) The original image with haze removed

Figure 4: The performance of ATM on a sample photo taken during day time.

(a) The original image (b) The haze component extracted by ATM (c) The original image with haze removed

Figure 5: The performance of ATM on a sample photo taken during night time.

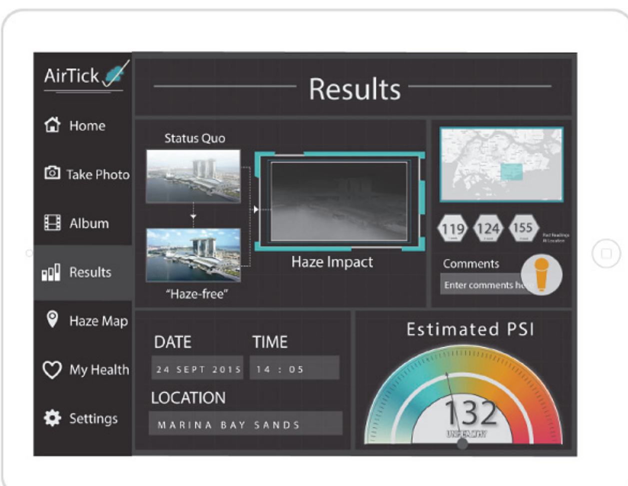

Figure 6: The AirTick prototype user interface.

the seriousness of haze in his/her surrounding environment. The contrast between the original image and the one with the haze component removed serves the purpose of priming the users with to aim to build awareness of the impact of haze and the importance of clean air. Additional information such as the time and location the air quality sensing action has taken place is also displayed in the user interface for easy reference. The estimated PSI value is displayed in the format of a meter dial colour coded to reflected the seriousness of health threat of the currently detected haze level based on NEA categorizations.

Experimental Evaluation

In order to train DNA and evaluate the performance of AirTick, we crawled 1,000 photos taken in Singapore between 2014 and 2016 from online sources. We use the creation date of the photo files as the date they were taken. Of these photos, 500 were taken during day time, and 500 were taken during night time. The selected night time photos are ones which include visible light sources (e.g., street lamps, headlights from vehicles). We ensured that the photos included in the dataset were taken after 1 April 2014 as the PSI value was computed differently before that date. Based on the creation date of these photos, we obtained the corresponding PSI values from the NEA website historical PSI data. In this way, a dataset has been constructed with photos and the corresponding PSI values (which are treated as the ground truth). We divide the dataset into a training set consisting of 800 and a testing dataset of 200 photo-PSI pairs.

The key metrics used to evaluate AirTick are:

1. PSI Estimation Accuracy:

$$acc = 1 - (|PSI - PSI_{est}|)/PSI \qquad (6)$$

where PSI is the ground truth PSI value, and PSI_{est} is the PSI value estimated by AirTick; and

2. Processing time: the time taken per pixel to compute the estimated PSI value.

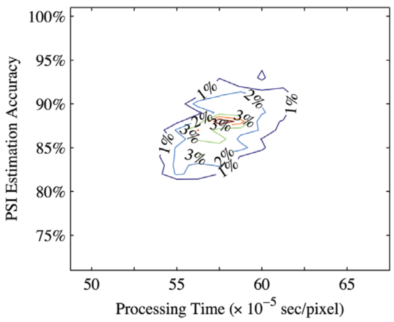

(a) Day time photos

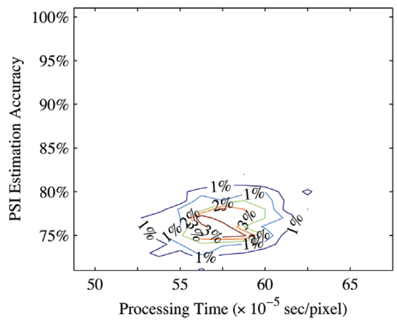

(b) Night time photos

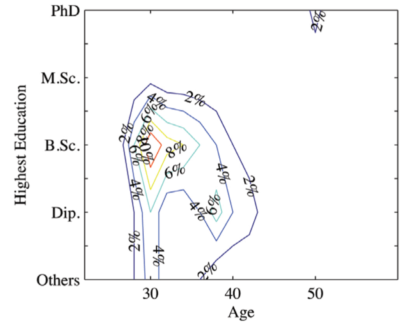

(c) Study participants' demographics

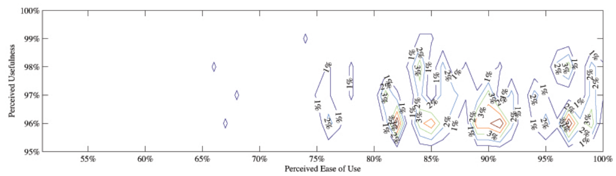

(d) Empirical study results

Figure 7: Evaluation results.

Figure 7(a) shows the performance of AirTick on photos taken during day time. The average accuracy value is 87%. Figure 7(b) shows the performance of AirTick on photos taken during night time. The average accuracy value is 75%, which is significantly less than the case of daytime photos. This is because, at night time, the airlight level is not as strong as during day time. The significant portions of dark areas in the photos make extracting haziness highly challenging. Currently, for night operation, AirTick requires the presence of visible light sources (e.g., street lamps, car headlights) in the photos taken in order to work. Nevertheless, the DNA component was able to compensate for night operation to some extent to achieve reasonably high estimation accuracy. There is no significant difference between the distributions of processing time for day time and night time photos. The average time required to estimate the PSI value of a given photo is around 5.8×10^{-6} seconds per pixel.

Empirical Evaluation

In order to gauge how users feel about AirTick, we conducted empirical user studies in collaboration with two Residents' Committees (RCs) in Singapore. In total, 100 local residents participated in the studies. The demographics of the study participants are summarized in Figure 7(c). The majority of them are in their early thirties with a university degree or a diploma. This group is considered technology savvy. Nevertheless, there are also some participants who are in the 40 to 50 age group. They were given the AirTick prototype system to evaluate and asked to comment on its perceived usefulness and ease-of-use through a Standard Usability Survey (SUS) (Bangor, Kortum, and Miller 2008).

The study participants feel strongly that AirTick is highly useful (with an average perceived usefulness score of 96%) (Figure 7(d)). The perceived ease-of-use is generally very high among younger participants (with an average perceived ease-of-use score of 82%). The participants in the 40 to 50 age group generally perceive AirTick to be less easy to use than the 30+ age group. Nevertheless, the perceived ease-of-use scores from these age groups are also reasonably high (with perceived ease-of-use scores all above 65%).

Discussion and Future Work

Being a free, convenient and scalable air quality monitoring tool powered by AI techniques, AirTick is ideally suited to improving public awareness of air quality. Although awareness by itself does not solve the problem, it fosters a sense of personal connection which can help convert many into environmentally conscience and responsible citizens, thus inducing more responsible and sustainable behaviours. AirTick can be a valuable smart governance tool for building sustainable and liveable cities.

At present, the accuracy achieved by AirTick has significant room for improvement, especially for night time operations. We will explore whether greater accuracy could be obtained by sharing information between several different AirTick users in close proximity to each other, both temporally and spatially. For this, we will leverage our previous works in spatial crowdsourcing (Miao et al. 2016) and collaborative crowdsourcing (Yu et al. 2015; Pan et al. 2016).

The psychological priming technique employed by AirTick has shown potential in enhancing user experience in the current study. We will further explore the impact of such techniques in terms of incentive mechanisms (Tao et al. 2011; Liu et al. 2015), trust (Yu, Shen, and An 2012; Mei et al. 2014; Liu et al. 2014; Yu et al. 2014), and influencing users' emotions (Lin et al. 2015).

In addition to air quality analytics, user contributed photos taken with AirTick will contain additional useful information. With the location, camera orientation, time and visible light source information associated with each photo, it is possible extract useful knowledge on the duration and strength of sunlight over diverse ranges of locations and time (Spitschan et al. 2016). This knowledge can be used to help urban planners optimize the deployment of renewable energy devices (e.g., solar panels) in the future.

Acknowledgements

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its IDM Futures Funding Initiative; and the Interdisciplinary Graduate School, Nanyang Technological University, Singapore; and the Lee Kuan Yew Post-Doctoral Fellowship Grant.

References

Anderson, J. O.; Thundiyil, J. G.; and Stolbach, A. 2012. Clearing the air: A review of the effects of particulate matter air pollution on human health. *Journal of Medical Toxicology* 8(2):166–175.

Bangor, A.; Kortum, P. T.; and Miller, J. T. 2008. An empirical evaluation of the system usability scale. *International Journal of Human-Computer Interaction* 24(6):574–594.

Chan, F. 2015. $47b? Indonesia counts costs of haze. *The Straits Times* http://www.straitstimes.com/asia/47b–indonesia–counts–costs–of–haze.

Hauck, H.; Berner, A.; Gomiscek, B.; Stopper, S.; Puxbaum, H.; Kundi, M.; and Preining, O. 2004. On the equivalence of gravimetric PM data with TEOM and beta-attenuation measurements. *Journal of Aerosol Science* 35(9):1135–1149.

Hinton, G. E., and Salakhutdinov, R. R. 2006. Reducing the dimensionality of data with neural networks. *Science* 313(5786):504–507.

Jutzeler, A.; Li, J. J.; and Faltings, B. 2014. A region-based model for estimating urban air pollution. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, 424–430.

Levin, A.; Lischinski, D.; and Weiss, Y. 2008. A closed form solution to natural image matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 30(2):228–242.

Lin, J.; Yu, H.; Miao, C.; and Shen, Z. 2015. An affective agent for studying composite emotions. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'15)*, 1947–1948.

Liu, Y.; Liu, S.; Fang, H.; Zhang, J.; Yu, H.; and Miao, C. 2014. Reprev: Mitigating the negative effects of misreported ratings. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, 3124–3125.

Liu, S.; Miao, C.; Liu, Y.; Yu, H.; Zhang, J.; and Leung, C. 2015. An incentive mechanism to elicit truthful opinions for crowdsourced multiple choice consensus tasks. In *Proceedings of the 2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'15)*, 96–103.

Liu, C.; Tsow, F.; Zou, Y.; and Tao, N. 2016. Particle pollution estimation based on image analysis. *PLoS ONE* 11(2):doi:10.1371/journal.pone.0145955.

Mei, J.-P.; Yu, H.; Liu, Y.; Shen, Z.; and Miao, C. 2014. A social trust model considering trustees' influence. In *Proceedings of the 17th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA'14)*, 357–364.

Meyer, D., and Schvaneveldt, R. 1971. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90(2):227–234.

Miao, C.; Yu, H.; Shen, Z.; and Leung, C. 2016. Balancing quality and budget considerations in mobile crowdsourcing. *Decision Support Systems* 90:56–64.

Muir, D., and Laxen, D. P. 1995. Black smoke as a surrogate for PM10 in health studies? *Atmospheric Environment* 29(8):959–962.

2014. PSI. *National Environment Agency (NEA), Singapore* http://www.nea.gov.sg/anti–pollution–radiation–protection/air–pollution–control/psi/psi.

Pan, Z.; Miao, C.; Tan, B. T. H.; Yu, H.; and Leung, C. 2015. The effects of familiarity design on the adoption of wellness games by the elderly. In *Proceedings of the 2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'15)*, 387–390.

Pan, Z.; Yu, H.; Miao, C.; and Leung, C. 2016. Efficient collaborative crowdsourcing. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*, 4248–4249.

Ruppecht, E.; Meyer, M.; and Patashnick, H. 1992. The tapered element oscillating microbalance as a tool for measuring ambient particulate concentrations in real time. *Journal of Aerosol Science* 23(1):635–638.

Silberzahn, R., and Uhlmann, E. L. 2015. Crowdsourced research: Many hands make tight work. *Nature* 526:189–191.

Soeriaatmadja, W. 2015. Thick haze grounds firefighting copters. *The Straits Times*.

Spitschan, M.; Aguirre, G. K.; Brainard, D. H.; and Sweeney, A. M. 2016. Variation of outdoor illumination as a function of solar elevation and light pollution. *Scientific Reports* 6(26756):doi:10.1038/srep26756.

Tao, X.; Shen, Z.; Miao, C.; Theng, Y.-L.; Miao, Y.; and Yu, H. 2011. Automated negotiation through a cooperative-competitive model. *Studies in Computational Intelligence* 319:161–178.

2014. 7 million premature deaths annually linked to air pollution. *WHO* http://www.who.int/mediacentre/news/releases/2014/air–pollution/en/.

Yu, H.; Shen, Z.; Miao, C.; An, B.; and Leung, C. 2014. Filtering trust opinions through reinforcement learning. *Decision Support Systems* 66:102–113.

Yu, H.; Miao, C.; Shen, Z.; Leung, C.; Chen, Y.; and Yang, Q. 2015. Efficient task sub-delegation for crowdsourcing. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, 1305–1311.

Yu, H.; Miao, C.; Leung, C.; Chen, Y.; Fauvel, S.; Lesser, V. R.; and Yang, Q. 2016. Mitigating herding in hierarchical crowdsourcing networks. *Scientific Reports* 6:doi:10.1038/s41598–016–0011–6.

Yu, H.; Shen, Z.; and An, B. 2012. An adaptive witness selection method for reputation-based trust models. In *Proceedings of the 15th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA'12)*, 184–198.

Yu, H.; Shen, Z.; and Leung, C. 2013. From internet of things to internet of agents. In *Proceedings of the 2013 IEEE International Conference on Internet of Things (iThings'13)*, 1054–1057.