INVITED PAPER

Human Sensor Networks for Improved Modeling of Natural Disasters

A human sensor network incorporated into geophysical models, together with satellite observations and sensor measurements, is proposed in this paper; oil spill predictions are used to validate improvements.

By Oleg Aulov and Milton Halem

ABSTRACT | In this paper, we present a novel approach that views social media (SM) data as a human sensor network. These data can serve as a low-cost augmentation to an observing system, which can be incorporated into geophysical models together with other scientific data such as satellite observations and sensor measurements. As a use case scenario, we analyze the Deepwater Horizon oil spill disaster. We gather SM data that mention sightings of oil from Flickr, geolocate them, and use them as boundary forcings in the General NOAA Oil Modeling Environment (GNOME) software for oil spill predictions. We show how SM data can be incorporated into the GNOME model to obtain improved estimates of the model parameters such as rates of oil spill, couplings between surface winds and ocean currents, diffusion coefficient, and other model parameters.

KEYWORDS | Data mining; human sensor networks; natural disasters; oil spill trajectory forecast; situational awareness; social media (SM)

Manuscript received June 8, 2011; revised February 16, 2012; accepted April 1, 2012. Date of publication July 10, 2012; date of current version September 14, 2012. This work was supported by the National Science Foundation (NSF) under the RAPID MRI award: An Interactive "Human Sensor Web" for Improved Model Predictions of the Dispersion of the Deepwater Horizon Gulf Oil Spill, Grant 106162 and the NSF I/UCRC Fundamental Research Supplement: Collaborative Research: An Interactive Situational Awareness Simulation—A View from the Clouds, Grant 1032040.
The authors are with the Department of Computer Science and Electrical Engineering, University of Maryland Baltimore County, Baltimore, MD 21250 USA (e-mail: aulov.oleg@gmail.com; halem@umbc.edu).

Digital Object Identifier: 10.1109/JPROC.2012.2195629

I. INTRODUCTION

People mainly use social media (SM) for social interaction. As a result of its popularity there is an abundance of SM data in various forms such as YouTube videos, Flickr images, tweets and blog posts that are mined for opinions, sentiment analysis, popularity ratings, and a variety of other purposes. Methods of extracting scientific data from SM and using it in physiological models are beginning to be researched [1], [2]. In this work, we describe an approach of viewing SM as a human sensor network (HSN), and information retrieval from SM as sensors in the field that observe a certain event and produce measurements of it. As a use case scenario, we look at the Deepwater Horizon oil spill disaster. In the aftermath of the disaster, the public was very active in discussing its impact and implications across a range of SM outlets. Many people who witnessed firsthand the damage caused by the oil—such as birds soaked in oil or tar balls washing up on the shore—reported their accounts in different SM outlets. People posted photos and videos of oiled beaches, tweeted from their smartphones when they were prohibited from swimming because of oil pollution, and so on. The National Incident Command under Admiral Thad Allen saw the potential of these data and utilized SM mining to gauge and monitor the mood of the public affected by the BP disaster [3]. Though the primary purpose of such online activity is social interaction between friends, increasingly traditional media outlets look to SM to improve their reporting and get hints about newsworthy events. In this paper, we describe how SM data can be used as physical observations to provide boundary forcing corrections to oil spill model predictions that employ generic parameterizations such as the coupling between the surface air and ocean drift velocities. These SM data can help adjust other

parameters in the oil spill model as well. We examined various SM outlets and collected the data that were related to the Gulf oil spill disaster. We converted these data for scientific use in geophysical models. For many of these data reports, we can extract observation location and temporal information. In addition, by comparing multiple reports from different observers (i.e., sensors) we can apply quality controls on the usefulness of the SM data. In Section II, we describe prior work that was done on the subject. We provide a summary of two prior research projects that used Twitter data—one to study earthquakes, and another to study air quality. We then describe the General NOAA Oil Modeling Environment (GNOME) software that is used by the Emergency Response Division (ERD) of NOAA's Office of Response and Restoration (OR&R) to predict the propagation of oil plumes on the ocean surface. Section III describes the theory behind our work and the governing equations. Here we explain the Eulerian–Lagrangian particle model that is used at the core of oil modeling software. Section IV describes our data sources and how we harvested the data and assimilated it into the GNOME model. Section V describes the implementation approach. We describe how we process the SM data and present the results of the integration experiments as a series of GNOME trajectory forecasts with different parameters that explore the parametric sensitivities of the model. Section VI discusses the results of our experiments and explains their significance. Section VII provides the summary of our work together with conclusions.

II. PRIOR WORK

In this section, we review prior work that was done both in the field of SM and in geophysics. In the SM area we demonstrate Twitter Earthquake Detector, #uksnow snow map mashup, and AirTwitter air quality alert system. In the geophysics area, we present general NOAA operational modeling environment (GNOME) that is used to forecast the movement of oil plumes on the surface of the ocean. Since our work is based on GNOME forecasts, we will discuss it in greater detail.

A. U.S. Geological Survey: Twitter Earthquake Detector (TED)

Many examples have been noted of individuals tweeting about experiencing an earthquake, even before media outlets received official reports of the event. Tweets about an earthquake experienced appear within minutes after the event, as opposed to official scientific reports that can take up to 20 min, depending on the location of the earthquake. In response to this trend, U.S. Geological Survey (USGS) started a Twitter Earthquake Detector (TED) [2]. TED is a system that gathers in real time tweets related to earthquakes, and processes them to provide geolocated areas where people felt shaking. It can potentially improve the earthquake response products as well as hazard information gathering and delivery and could therefore improve the effectiveness of emergency response efforts. Fig. 1 shows a snapshot of TED in action displaying a map of the San Francisco, CA, area. Different colors of circles indicate frequency of tweets in that area that include the word "earthquake."

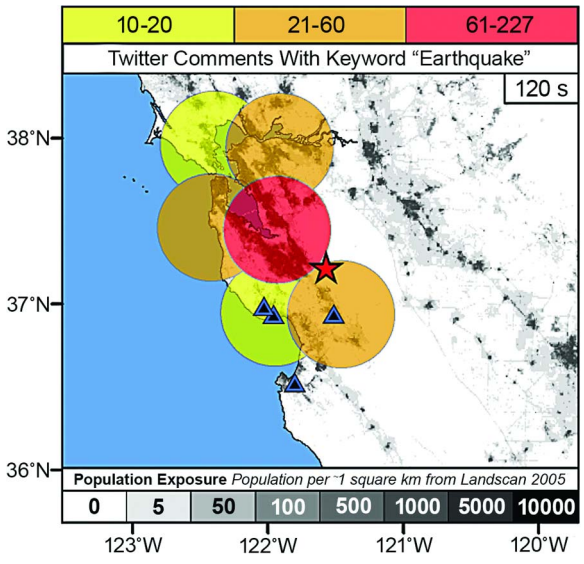

Fig. 1. *USGS Twitter Earthquake Detector (TED) displays the map of the San Francisco, CA, area during an earthquake event. Red star indicates the epicenter of the earthquake. Darker map regions indicate more densely populated areas. Different colors of circles indicate frequency of geolocated tweets that include the word "earthquake."*

During the Tohōku earthquake and tsunami of 2011 we conducted a preliminary investigation of the timeliness of SM outlets. The earthquake occurred at 14:46 JST (05:46 UTC) on March 11th, 2011. We discovered that for this particular event, Twitter was flooded with relevant tweets immediately, and Flickr had postings within about 40 min. Fig. 2 indicates the timeline of frequency of tweets with keywords "earthquake" and "Japan" the day of the earthquake. Note that there were some tweets prior to the event. Those were referencing relatively minor earthquakes that happened several days prior to the devastating one.

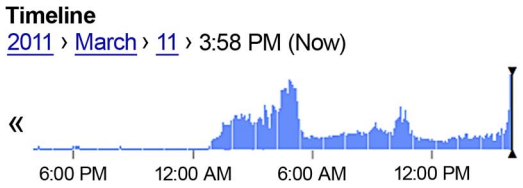

Fig. 2. *Timeline of tweets that include words "earthquake" and "Japan."*

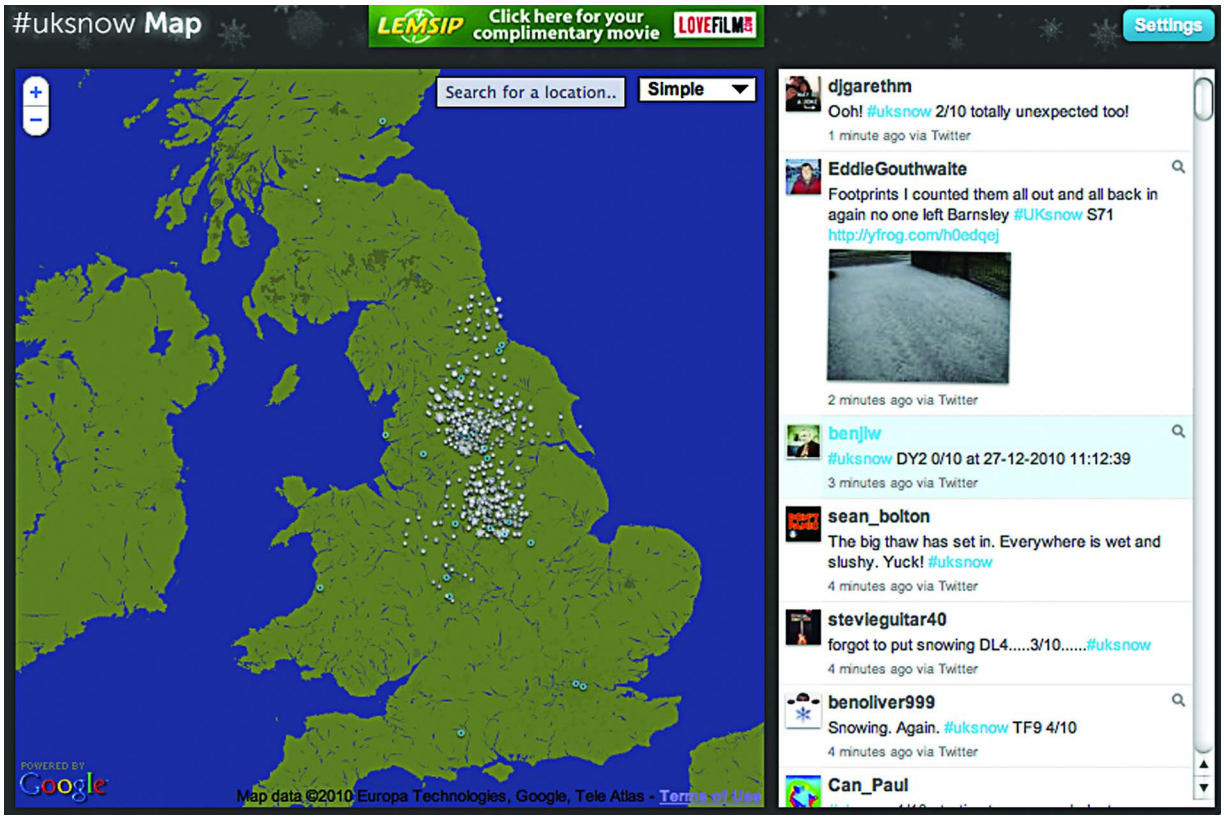

Fig. 3. *#uksnow Twitter mashup. Google Map of the United Kingdom on the left with white circles of varying sizes indicating snow conditions. Tweet feed on the right listing tweets that were used to generate the map on the left. Participants tweet about current weather conditions in their area rating the snow from 0 to 10, where 0 means no snow, 1–3 means a few snowflakes, 5 means steady snow, and 10 is the worst blizzard imaginable.*

B. AirTwitter

In the past decade, there has been a trend online to create mashups. Mashups are web-based applications that combine together or heavily rely on multiple other web-based applications to create a new type of application. AirTwitter is one such mashup application that monitors SM data, identifies air-quality- and pollution-related events, and records and monitors this information. AirTwitter aggregates RSS feeds from different SM sites such as Twitter, YouTube, Flickr, Delicious, and others. It processes the data to weed out all the feeds unrelated to air quality. Then, it establishes a baseline of normal frequency of air-quality data. As a result, AirTwitter provides an aggregated, preprocessed, single feed of data with the ability to detect air-quality-related events, such as volcano eruptions and forest fires [4].

C. #uksnow Map

This Google Maps/Twitter mashup online application, called #uksnow Map, tracks in real time snow reports and displays them on the map. The pound sign in the beginning of the name is a pun on the twitter hash tag used for this application. #uksnow Map is based on crowdsourcing. In this mashup, Twitter users are asked to report about the snow conditions in their area. Tweets should include the #uksnow hashtag, location, and snow rating on a scale from 0 to 10. The location is represented as a postal code, a town name, or a Twitter geotag (latitude and longitude). The snow scale is very loose—the snow is rated as 0 for no snow at all, 1–2 for a few flakes, 5 for a steady snow, and 10 for a blizzard. Attaching photos and including the depth of the snow is also encouraged. The application keeps track of tweets tagged with the #uksnow tag and displays them on a Google Maps map in real time. Fig. 3 demonstrates a screenshot of #uksnow Map. On the left portion of the screen is a Google Maps window zoomed into the United Kingdom. Superimposed white circles of varying diameter indicate the intensity of the snow. On the right section of the screen is a live tweet feed that is used to generate the map in real time [1].

D. General NOAA Operational Modeling Environment (GNOME)

GNOME is the NOAA forecast model that implements an oil spill trajectory prediction. It is used by the ERD of NOAA's OR&R during an oil spill and was used

for issuing daily operational oil spill forecasts during the Deepwater Horizon disaster. GNOME software uses surface winds, surface ocean currents, and other processes as input to predict the movement and spread of oil on the ocean surface as well as how the predicted oil trajectories might reach coastal beaches and islands. The model is affected by uncertainty in currents and winds observations and forecasts as well as in the rate of oil spill and the thickness of oil reaching the surface of the ocean. GNOME also has the capability of predicting the weathering behavior of pollutants. The model is initialized by setting up a spill scenario that takes as an input the shoreline, surface winds, surface ocean currents/tides, and the observed location of oil plumes on the surface of the ocean. As an output, GNOME produces an animation of how the plume moves, mixes, and weathers over time as well as a batch file with a series of data points representing time series of predicted locations of oil particles.

In its core, GNOME software uses basic Lagrangian–Eulerian particle tracking algorithms. The area of interest is divided into a grid and a certain arbitrary amount of oil in each cell is viewed as a single granule [Lagrangian element (LE)] that is influenced by the velocities of the universal movers such as winds and currents. The output of the GNOME model represents the LE of the spilled oil trajectory as splots (spill dots). Black splots represent the best guess trajectory estimate and red splots represent the minimum regret. Best guess trajectories are calculated under the assumption that the winds and currents data accurately represent the actual winds and currents over the period of the scenario, and that the initial input of splots representing the observed location of the sheen of oil from satellite updates are accurate as well. The best guess forecast trajectory takes into account the turbulence that is inherent in the surface winds and ocean surface currents. Red splots, on the other hand, represent the minimum regret trajectory. Minimum regret trajectory takes into account the inherent uncertainty of the winds and ocean currents models and specifies the boundary beyond which there is a high probability that the oil will not propagate (probability in the order of 90%). Minimum regret trajectories are useful for purposes of indicating trajectory uncertainties that are less probable than the best guess trajectory, but that may be potentially more destructive [5]–[7]. Fig. 4 demonstrates the GNOME model initialized with the coastline of the Gulf of Mexico. Purple arrows visualize the surface currents. The red and black areas are splots of minimum regret and best guess forecasts.

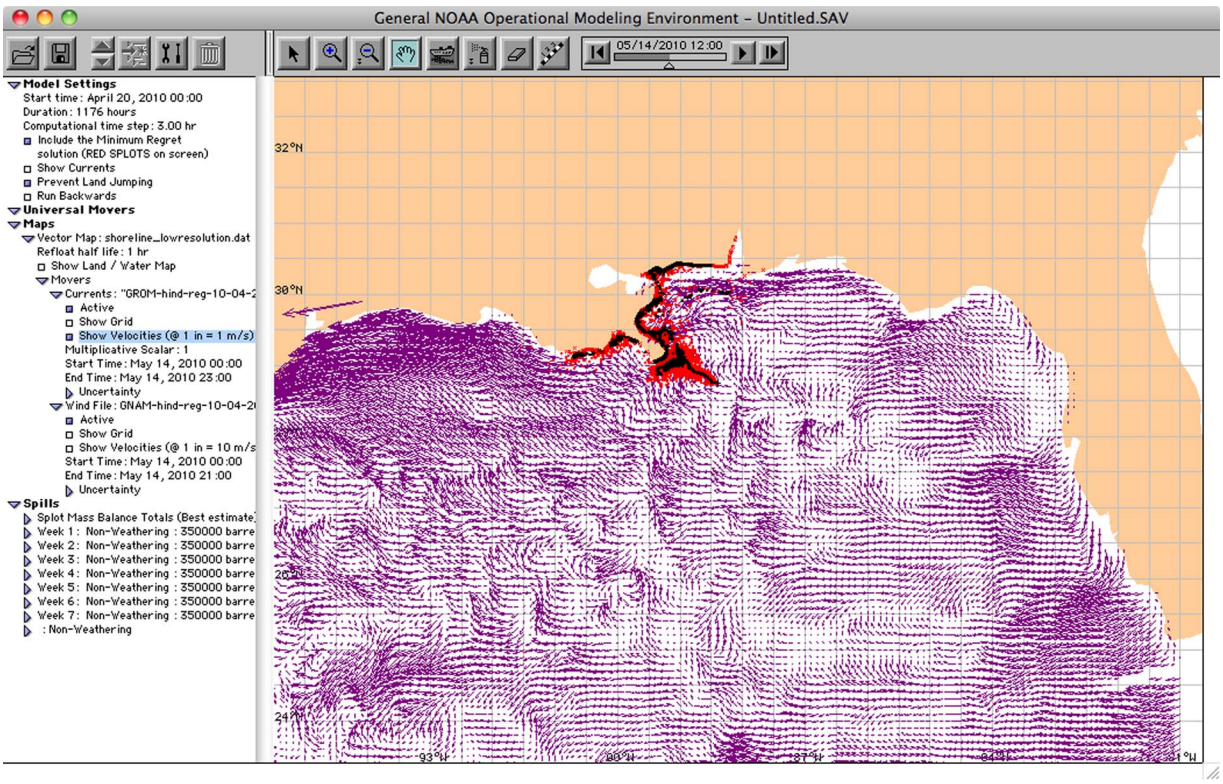

Fig. 4. *Screenshot of GNOME software initialized with surface winds, ocean currents, and coastline data for Deepwater Horizon oil spill scenario. Purple arrows indicate velocity vectors of surface currents. Black dots are LEs that indicate the best guess of the oil location while red dots are LEs that indicate the minimum regret of the oil location.*

GNOME has a few variables that are adjustable by the user and are generally set experimentally. Those variables include number of splots per spill, windage percentage range for each spill, pollutant release rate, along- and cross-current uncertainty percentage, wind speed scale, and total wind angle scale.

III. GOVERNING MODEL EQUATIONS FOR OIL SPILL TRAJECTORY DYNAMICS

The most common approach in forecasting floating objects on the water surface, such as oil plumes on the ocean surface, is using an implementation of basic Eulerian–Lagrangian equations. We assume that the thickness of the oil slick on the surface of the ocean is negligible in comparison to the thickness of the water [8]–[10]. The governing equations for the movement of the oil slick in the 3-D space x, y, z over time t are

$$\frac{\partial C_s}{\partial t} + \frac{\partial}{\partial x}(u_S C_s) + \frac{\partial}{\partial y}(v_S C_s) - \frac{\partial}{\partial z}\left(K_x \frac{\partial C_s}{\partial x}\right) + \frac{\partial}{\partial y}\left(K_y \frac{\partial C_s}{\partial y}\right) + \left(\alpha_1 v_b C_v + K_z \frac{\partial C_s}{\partial z}\right) - \gamma C_s - S_e - S_d + M_s(x,y) - D_s(x,y). \quad (1)$$

C_s is the concentration of oil on the surface of water; C_v is the volume of oil concentration in the suspended layer per volume of water; u, v, w are components of surface current velocity; α_1 is the coefficient of the probability of an oil droplet reaching the surface; γ is the surface oil dispersion coefficient; u_s, v_s are the components of the drift velocity; K is the diffusion coefficient in the respective direction; v_b is the buoyant velocity of the suspended oil parcels; S_e is the rate of evaporation; S_d is the rate of dissolution; M_s is the effect on the distribution of the surface oil due to the mechanical spreading; D_s is the effect on the distribution of surface oil due to the shoreline deposition. [8] According to Reddy, advection is the main mechanism that governs the drifting of the suspended oil and the surface oil slick. The drift velocity u_s, v_s is considered a weighed combination of the velocity of the surface currents with the velocity of the surface winds. The weighting parameters that are generally used to combine the air surface winds and the ocean currents are [8]

$$0.03 v_a + 1.1 v_c = v_s. \quad (2)$$

IV. DATA SOURCES

In this section, we describe the geophysical data that we used to initialize the GNOME model as well as the method of extracting the SM data from Flickr and converting it to geophysical data.

A. Shoreline Data

In order to make the oil spill movement forecasts, the GNOME model needs to be properly set up with the shoreline data that defines the map separating land from water, as well as the surface winds and ocean currents forecast data. We generated a custom shoreline map data file for the purpose of the boundary condition that specifies where the separation of ocean surface and land is, so that the model knows where the oil plume is floating and where it makes a landfall. The GNOME model expects a map file that contains a list of latitude/longitude points that represent a polygon of land. We extracted the USGS shoreline data of the Gulf of Mexico area from the Coastline Database hosted at NOAA's National Geophysical Data Center. [11] We used NOAA/NOS medium resolution coastline data designed for 1 : 70 000 scales. This coastline data consist of many lists of latitude/longitude tuples. Each list contains various numbers of those tuples and represents a small portion of the shoreline in the cylindrical–equidistant projection. Then, we linked all these lists together in the right order to create a single list that represented the entire shoreline of the Gulf coast. Since the model expects a polygon representing the land, and assumes that everything else is water, it was necessary for us to add arbitrary latitude/longitude points in the area of the state of New Mexico as well as the South Pacific area of Mexico. As a result, we got a polygon shape that represented the shape of the land around the Gulf of Mexico and that can be ingested into the GNOME model. Although in our model scenario the Pacific Ocean was abruptly starting after New Mexico, it did not matter since the forecasts and computations were strictly limited to a narrow region of the Gulf Coast and the area around the Macondo well. We also did not take into account in our model the unlikely possibility of oil being moved to Europe with the Gulf Stream and therefore we did not incorporate any shoreline data for European countries.

B. Ocean Currents and Surface Winds Data

We obtained the wind fields from NCDC that were generated by the NCEP ETA Regional Forecast model and the ocean current data from the ROM ocean model. The surface winds and ocean surface currents data were retrieved from the repository of the Department of Oceanography at Texas A&M University.

The atmospheric surface wind data are produced by NOAA's NCEP ETA-12 model and provide 24-h forecasts with output every 3 h on a regular grid with a grid spacing of 12 km. The ETA model, developed by Janjić and Mesinger, derives its name from the name of the model's vertical step mountain coordinate. The basis of the model is to minimize errors due to the gradient force

computation, advection, and diffusion. The vertical coordinate is defined by

$$\eta = \left[\frac{(p - p_t)}{(p_s - p_t)}\right] \times \left[\frac{(p_{\text{ref}}(Z_s) - p_t)}{(p_{\text{ref}}(0) - p_t)}\right] \quad (3)$$

where p_t is the pressure at the top of the domain, p_s is the pressure of the model's lower boundary, z_s is the elevation of the model's lower boundary, and p_{ref} is the reference pressure state [12]. For the purpose of our research, we are only interested in the first layer of the model that correlates to the winds on the ocean surface.

The ocean surface currents data are produced by the Regional Oceanographic Modeling System (ROMS) generated by the Texas General Land Office (TGLO). ROMS provides a 24-h hourly forecast, four times a day, on a regular grid. ROMS is a high-resolution, free-surface, hydrostatic, primitive-equation ocean model that uses terrain-following coordinates in the vertical curvilinear coordinates in the horizontal plane [13]. In its core, it is based on the S-coordinate Rutgers University Ocean Model (SCRUM) [14]. The primitive horizontal momentum equations of this model are given as

$$\frac{\partial u}{\partial t} + v\nabla u - fv = -\frac{\partial \phi}{\partial x} + \frac{\partial}{\partial z}\left(K_M \frac{\partial u}{\partial z}\right) + D_u + F_u \quad (4)$$

and

$$\frac{\partial v}{\partial t} + v\nabla v - fv = -\frac{\partial \phi}{\partial y} + \frac{\partial}{\partial z}\left(K_M \frac{\partial v}{\partial z}\right) + D_v + F_v \quad (5)$$

where f is the Coriolis parameter, D is the horizontal viscous and diffusive term, F is the forcing term, K_M is the vertical eddy viscosity, and φ is the dynamic pressure.

C. Social Media Model Input Data

For the purpose of this study, we collected data from Flickr. Flickr is a widely used image and video hosting site as well as a web services framework. It features many social networking traits such as the ability for users to add people to a list of their contacts, forming communities, tagging people on photos and videos, tagging content with keywords, the ability to comment on photos, etc. As a result, Flickr is not only a popular SM portal in its own right, but also widely used by bloggers to host images that are imbedded in their blogs. Unless explicitly disabled by the user, photos and videos posted on Flickr include EXIF metadata such as date and time when the photo was taken, camera make and model, camera settings, and geolocation. Although the ratio of geolocated Flickr images is very small, we expect it to grow rapidly due to high-tech manufacturers objective to embed Global Positioning System (GPS) devices not only in smartphones, but also in regular photo cameras [15]. Currently, the vast majority of geolocated photos on Flickr are taken with a GPS-enabled smartphone such as iPhone, Blackberry, or Android-based device. If the EXIF metadata is not available, the user has an option of geotagging the photo by hand by selecting the location on the map where the photo was taken; if the timestamp is not available, Flickr will automatically assume that the upload time is the time the photo was taken. We started our data mining task with a simple search on the Flickr website for the query "BP oil spill," which returned over 20 000 results. Many of those images were related to protests against BP, political events related to the disaster, and other related events that were of no practical use for us, since we were only looking for oil HSN data in the form of images that evidenced oil slicks on the water and oil tar balls washing up on the shores. For developers, Flickr provides an application program interface (API) access to their service that allows other software to directly query and interact with Flickr resources. For our work, Flickr API proved to be much more flexible and powerful than the services accessible via the website. Using Flickr API, we were able to request only the images that were geotagged as well as supply our search query with a bounding box of lower left corner at 28.07198, −95.668945 and upper right corner at 31.203405, −85.825195 to retrieve only the images from the Gulf coast area. We executed two search queries—one for "tar balls" and the other one for "oil spill"—during the period of April 20, 2010 to October 20, 2010. They resulted in two disjoint sets of 22 and 168 images, respectively, for a total of 190 images. Fig. 5 shows an example of REST API query, and Fig. 6 shows the corresponding response from the Flickr API. Note attributes

```
http://api.flickr.com/services/rest/?method=flickr.photos.search&api_key=70920bca
63b7452f4ff6a7bdbb7f3f75&tags=tar+balls&min_taken_date=2010-04-20+00%3A00%3A00
&max_taken_date=2010-10-20+00%3A00%3A00&bbox=-95.668945%2C+28.07198%2C+-85.825195
%2C+31.203405&has_geo=1&extras=geo%2C+path_alias%2C+date_taken&auth_token=
72157626247331502-22c36f6dd37efa88&api_sig=89078181e2169738d46489f44f070843
```

Fig. 5. *Example of a Flickr API query.*

```
<photo id="5016704044" owner="37281343@N03" secret="9db693ebb5" server="4144"
farm="5" title="DSC_1225" ispublic="1" isfriend="0" isfamily="0" latitude=
"30.371133" longitude="-86.918726" accuracy="16" place_id="uwvhGpebBZlHgiRz"
woeid="2457354" geo_is_family="0" geo_is_friend="0" geo_is_contact="0"
geo_is_public="1" pathalias="mmmeeks" datetaken="2010-08-04 17:48:24"
datetakengranularity="0" />
```

Fig. 6. *Example of a Flickr API response.*

latitude," "longitude," "accuracy," and "date taken." Accuracy is a value in the range from 1 to 16 that represents the accuracy of the geolocation with the world level being 1, country level being about 3, region about 6, city being 11, and street being 16 (default is 16) [16]. We wrote a script that parses the XML response received from Flickr API and that generates a map with desired markers corresponding to the latitude/longitude tuple of each image. Using this script, we generated a series of maps of the Gulf coast of Mexico and incrementally superimposed Flickr data that we collected in the previous step at the corresponding latitude/longitude coordinates. As a result we got a series of maps indicating the sightings of oil or tar balls along the shoreline of the coast at different times. Fig. 7 displays the maps that were generated for different periods. Fig. 7(a) shows the points for the period from April 21 through April 27. Here we only see points around the location of the Deepwater Horizon drilling rig. Fig. 7(b) shows points for the period of April 21 through May 11, and here we observe additional points for the sightings of oil washing up in the area of Fort Morgan and Fort McRee. Fig. 7(c) is for the period from April 21 through May 30, and here we see additional points in the regions of Black Bay and Ship Island. Fig. 7(d) is for the period of April 21 through September 9 and we observe that many more areas in-between start to fill in.

V. APPROACH

In this section, we describe how we process the SM data and present the results of the integration experiments as series of the GNOME trajectory forecasts with different parameters that explore the parametric sensitivities of the model.

A. GNOME Parameter Sensitivity Experiments

Once the GNOME model was initialized for the area of the Deepwater Horizon oil spill, we introduced the spill itself at the exact location of the rig and for simplicity assumed that all the oil was spilling on the surface. In our first experiment, we altered the rate of spill and the number of splots used to represent the oil. We ran the scenario from the moment that the spill started until June 8. Fig. 8 summarizes our results. Fig. 8(a) indicates a spill at a rate of 7000 barrels a week with seven barrels per splot, Fig. 8(b) indicates a spill at a rate of 35 000 barrels a week with 35 barrels per splot, Fig. 8(c) indicates a spill at a rate of 35 000 barrels a week with seven barrels per splot, and Fig. 8(d) indicates a spill at a rate of 350 000 barrels a week with seven barrels per splot. Fig. 8(a) and (b) allows us to analyze the influence of the spill rate on the oil propagation given a constant number of splots used in the model. The rate changes from 7000 barrels a week to 35 000 barrels a week, but the number of splots is kept at 1000 per week. Fig. 8(b) and (c) allows us to analyze the influence of the number of splots on the oil propagation given a constant spill rate used in the model. The rate of spill is kept at 35 000 per week and the number of splots changes from 1000 per week to 5000 per week. Fig. 8(d) increases both the rate of spill and the number of splots in order to have a bigger picture of the influence of those two variables on the GNOME model. As a result of this experiment, we observed from Fig. 8(a) and (b) that the amount of oil released does not alter the direction of the oil flow or the ultimate location where the model predicts it will make landfall. Likewise, the number of splots does not alter the direction of the oil flow or the ultimate location where the model predicts it will make landfall, as can be seen in Fig. 8(b) and (c). Fig. 8(d) indicates the worst case scenario (highest rate of spill) with a substantially higher number of splots. The resulting output had exactly the same pattern of oil; however, it had a much better resolution (less grainy). It is important to mention that the model run that generated Fig. 8(d) also took a substantially longer time to run.

B. Assimilation of Social Media Data With GNOME Model

Now that we developed and presented a method to aggregate SM data, converted it into a format of latitude, longitude, timestamp triplets, and plotted it on the map, we were able to combine the HSN SM data with different forecasts of the GNOME model that correlate with different variable parameters of the model. We picked June 8 as our comparison date and ran the GNOME model from the day the spill started until June 8. For the first experiment, our spill was set up at a rate of 350 000 barrels per week with a representation of 1000 LEs per week. Then, we compared the results of different runs with the HSN data and summarized it in Figs. 9 and 10. Green stars indicate the SM data from Flickr, while red dots indicate the GNOME forecast of the oil spill. Fig. 9(a) shows the correlation of the default setting of the GNOME model with

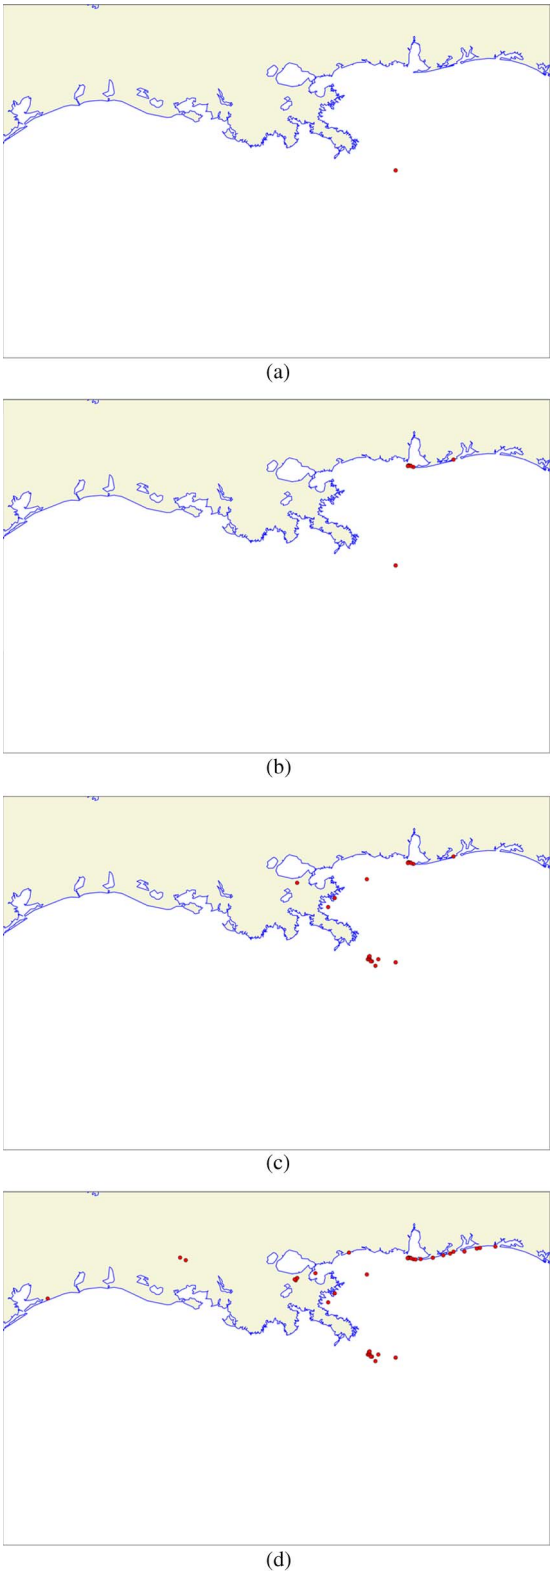

Fig. 7. *Accumulative Flicker data superimposed on a regional map. We can observe how as time goes by more and more areas get covered with oil: (a) April 21, 2010–April 27, 2010; (b) April 21, 2010–May 11, 2010; (c) April 21, 2010–May 30, 2010; and (d) April 21, 2010–September 9, 2010.*

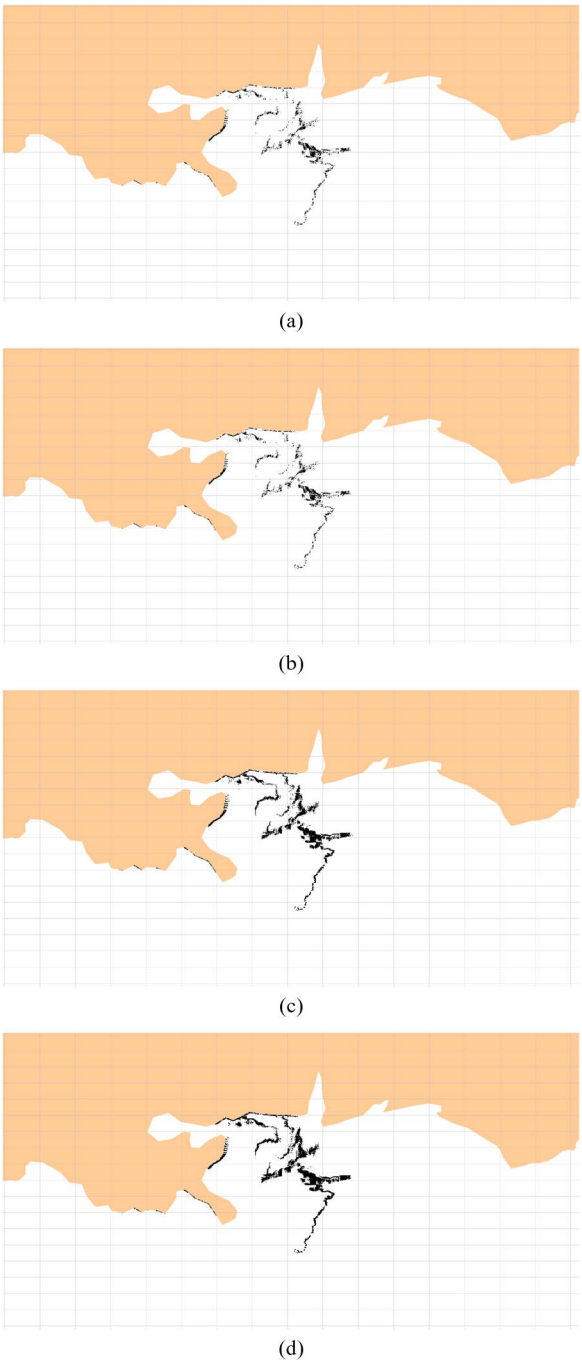

Fig. 8. *The results of different model runs show that the amount of oil released and the number of splots used do not alter the direction of the oil flow and the ultimate location where the model predicts it will make a landfall: (a) and (b) show that increase in the amount of oil does not alter the trajectory; (b) and (c) show that increase in the number of splots (LEs) does not alter the trajectory; (d) has increase in both the amount of the pollutant and the number of splots. The resulting figure clearly has a higher resolution of the predicted location of the pollutant, however the trajectory stays exactly the same. (a) 7000 barrels and 1000 LEs per week, seven barrels per splot; (b) 35 000 barrels and 1000 LEs per week, 35 barrels per splot; (c) 35 000 barrels and 5000 LEs per week, seven barrels per splot.*

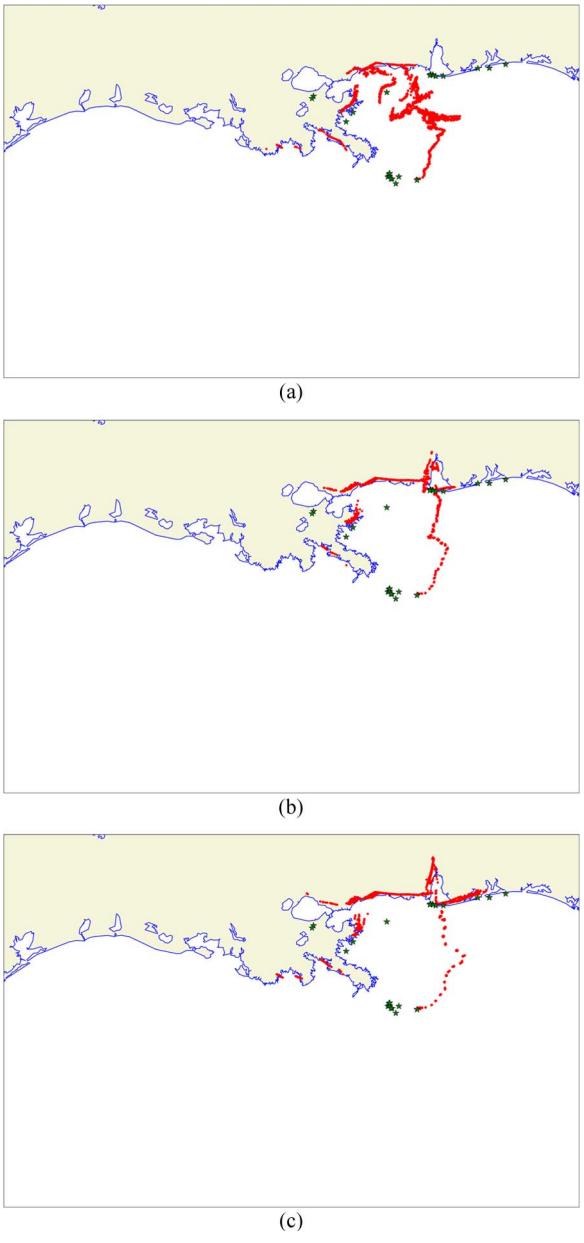

Fig. 9. *Comparison of the results of different GNOME model trajectory forecasts and their correlation to SM data. Red dots indicate the LEs from GNOME model while green stars represent SM data. We can see that depending on the settings of the model parameters some forecasts correlate better with SM data than others: (a) April 20–June 8, 2010 trajectory forecast —350 000 barrels per week, 1000 LEs per week with 1%–4% windage and no diffusion; (b) April 20–June 8, 2010 trajectory forecast—350 000 barrels per week, 1000 LEs per week with 5%–8% windage and no diffusion; and (c) April 20–June 8, 2010 trajectory forecast—350 000 barrels per week, 1000 LEs per week with 9%–12% windage and no diffusion.*

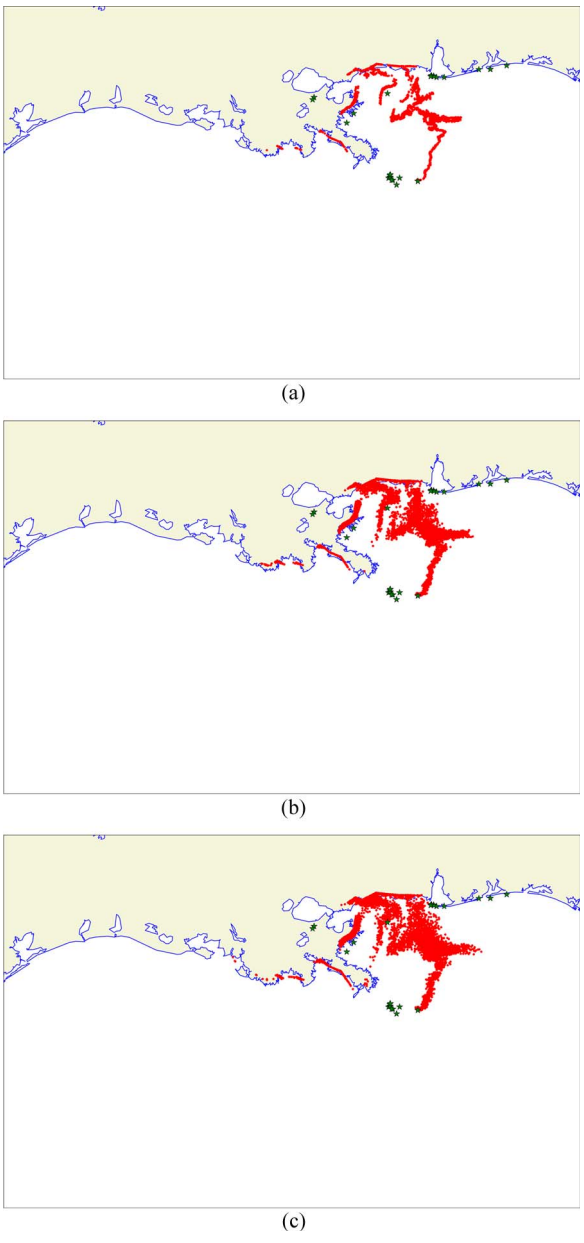

Fig. 10. *Comparison of the results of different GNOME model trajectory forecasts and their correlation to SM data. Red dots indicate the LEs from GNOME model while green stars represent SM data. We can see that depending on the settings of the model parameters some forecasts correlate better with SM data then others: (a) April 20–June 8, 2010 trajectory forecast—350 000 barrels per week, 1000 LEs per week with 1%–4% windage; (b) April 20–June 8, 2010 trajectory forecast—350 000 barrels per week, 1000 LEs per week with 1%–4% windage, and diffusion coefficient of 100 000 cm^2/s with uncertainty factor of 2; and (c) April 20–June 8, 2010 trajectory forecast—350 000 barrels per week, 1000 LEs per week with 1%–4% winds and diffusion coefficient of 200 000 cm^2/s with uncertainty factor of 2.*

wind factor of 1%–4%. Fig. 9(b) shows the correlation of HSN with the GNOME forecast set to wind factor of 5%–8%, and Fig. 9(c) is for wind factor of 9%–12%. For the second experiment, we kept the same date of June 8 with 350 000 barrels spilling per week with the representation of 1000 LEs per week and windage of 1%–4%. This time we introduced diffusion into the model and altered the

Table 1 RMS Error Results

Parameters	RMS Error
Windage: 1-4 %	20.443
Windage: 5-8 %	1.365
Windage: 9-12%	11.499
Diffusion: 100,000 cm2/sec	12.096
Diffusion: 200,000 cm2/sec	5.931

This table summarizes the RMS errors between social media data and different runs of the GNOME model.

diffusion variables. Fig. 10 displays the results of this experiment on a map. Fig. 10(a) shows no diffusion. Fig. 10(b) shows the forecast with the diffusion coefficient of 100 000 cm^2/s and an uncertainty factor of 2. Fig. 10(c) shows the forecast with a diffusion coefficient of 200 000 cm^2/s and an uncertainty factor of 2. Now that we have run multiple experiments and have gotten both SM data and GNOME forecast data in the same format, we can analyze the results of our experiments.

VI. RESULTS

We assumed that the SM data were the ground truth and we compared GNOME model forecasts to that ground truth by calculating the root mean square (RMS) error. The RMS calculation was performed as follows: for each point of SM data, we found the closest LE point from the forecast and we calculated the geometric distance in kilometers. We squared each such distance, summed all the distances together, and extracted the square root of that sum. Our calculations are summarized in Table 1. The first three entries are for experiments with different winds and currents combinations and no diffusion. The last two entries assume the windage of 1%–4% and introduce the comparison of two different diffusion coefficients.

It is important to point out that a change of windage of a few percentage points (from 1%–4% to 5%–8%) resulted in the RMS change of an order of magnitude (from 20.4 to 1.3). Similar outcome is observed when diffusion is introduced. We have also conducted additional studies—not presented in this paper—on the size of the initial oil plume, in particular, a simulated hand-drawn map of the

Fig. 11. *Screenshot of the Environmental Response Management Application (ERMA) Gulf Response mapping site displaying MODIS image of sun-glint reflecting from the oil plume of the Deepwater Horizon disaster on May 24, 2010 as well as NESDIS anomaly analysis composite product derived from COSMO SKYMED-2, MODIS TERRA, and RADARSAT-2 for the same date (green superimposed polygon) [17].*

oil spill similar to that of the sun glint reflection observed by the moderate-resolution imaging spectroradiometer (MODIS) instrument on the Terra satellite taken on May 24, 2010, shown in Fig. 11. This study showed increased beached oil trajectories, indicating the need for realistic quantitative mapping of the thickness of inferred oil spill images [18]. Currently, the National Environmental Satellite Data and Information Service (NESDIS) provides a composite product that is created by hand by their trained technicians. Some research has been conducted on automating such processes using machine-learning techniques, such as the work of Corucci on oil spill classification of multispectral satellite images using neurofuzzy technique [19] and the work of Lary on dust source classification from satellite imagery using self-organizing maps[20]. We are actively collaborating with both authors in order to bring those approaches to maturity.

VII. SUMMARY

We processed the SM data and converted it to physical observations that list latitude, longitude, and timestamp when the oil landfall was observed. The latitude and longitude can be obtained in different ways depending on the source of the post. In the case of a tweet that was posted from a smartphone, this information is available in the metadata of the tweet itself since Twitter geolocates tweets that are posted from smartphones. In the case of Flickr images, some cameras, especially those built into smartphones, often automatically geotag the photos, and this information is often preserved in the metadata of the image. In the cases where the tweet was sent from a computer, we can still get a more coarse location from the geolocation of the IP address of the Twitter user. Twitter provides such a coarse location as well. In the case where the photograph or the video was not geotagged, as well as in the case of most blog posts, we will have to extract such information from the textual content such as image descriptions or reader comments. Such extraction can be automated using text analysis tools and named entity recognizers. In our case, we used only the data that were already geotagged. We processed SM data from Flickr to be in the format of the observational geophysical data, and used it as a boundary condition to assess the sensitivity and agreement of time-dependent parameters in the GNOME model with SM data. We quantified the differences between the forecast and the SM observations by calculating the RMS error. We observed that minor changes in initial conditions of the forecast model can lead to an order of magnitude increase in consistency with specified Flickr data. ■

Acknowledgment

The authors would like to thank the Emergency Response Division (ERD) of NOAA's Office of Response and Restoration (OR&R) for technical support in utilizing GNOME oil spill modeling software.

REFERENCES

[1] C. Beaumont, "Google Maps and Twitter mashup show UK snowfall," *The Telegraph*, Feb. 2, 2009. [Online]. Available: http://www.telegraph.co.uk/technology/4433778/Google-Maps-and-Twitter-mashup-shows-UK-snowfall.html.

[2] U.S. Department of the Interior Recovery Activities, *U.S. Geological Survey: Twitter Earthquake Detector (TED)*. [Online]. Available: http://recovery.doi.gov/press/us-geological-survey-twitter-earthquake-detector-ted/

[3] J. Achenbach, "An excerpt from Joel Achenbach's 'A Hole at the Bottom of the Sea'," *The Washington Post*, Apr. 3, 2011. [Online]. Available: http://www.washingtonpost.com/lifestyle/style/an-excerpt-from-joel-achenbachs-a-hole-at-the-bottom-of-the-sea/2011/03/30/AFNy6bXC_story.html

[4] Federation of Earth Science Information Partners, *Air Quality Work Group*. [Online]. Available: http://wiki.esipfed.org/index.php/Air_Quality_Work_Group

[5] C. Beegle-Krause, "General NOAA Oil Modeling Environment (GNOME): A new spill trajectory model," in *Proc. Int. Oil Spill Conf.*, 2001. [Online]. Available: http://www.iosc.org/papers_posters/00891.pdf

[6] H. M. R. D. Office of Response and Restoration, *GNOME (General NOAA Oil Modeling Environment) Users Manual*, National Oceanic and Atmospheric Administration (NOAA), Tech. Rep., 2002.

[7] C. Beegle-Krause, "GNOME: NOAA's next-generation spill trajectory model," in *Proc. MTS/IEEE OCEANS Conf.*, 1999, vol. 3, pp. 1262–1266.

[8] G. Reddy and M. Brunet, "Numerical prediction of oil slick movement in Gabes Estuary," Transoft International, EPINAY/SEINE Cedex, France, 1997.

[9] P. D. Yapa, "Oil spill processes and model development," *J. Adv. Mar. Technol.*, vol. 11, pp. 1–22, 1994.

[10] P. Yapa, "Modelling river oil spills: A review," *J. Hydraulic Res.*, vol. 32, pp. 765–782, 1994.

[11] R. Signell, *Coastline Extractor*. [Online]. Available: http://www.ngdc.noaa.gov/mgg/coast/

[12] T. L. Black, "The new NMC mesoscale eta model: Description and forecast examples," *Weather Forecast.*, vol. 8, pp. 265–278, Jun. 1994.

[13] A. Shchepetkin and J. McWilliams, "The regional oceanic modeling system (ROMS): A split-explicit, free-surface, topography-following-coordinate oceanic model," *Ocean Model.*, vol. 9, no. 4, pp. 347–404, 2005.

[14] Y. Song and D. B. Haidvogel, "A semi-implicit ocean circulation model using a generalized topography-following coordinate system," *J. Comput. Phys.*, vol. 1, no. 115, pp. 228–244, 1994.

[15] R. Fairlie, "Samsung Compact Camera Adds GPS," *The New York Times*, Jan. 22, 2010. [Online]. Available: http://gadgetwise.blogs.nytimes.com/2010/01/22/samsung-compact-camera-adds-gps/.

[16] Flickr, *The App Garden: flickr.photos.geo.setLocation*. [Online]. Available: http://www.flickr.com/services/api/flickr.photos.geo.setLocation.html

[17] Environmental Response Management Application (ERMA), *ERMA Gulf Response*. [Online]. Available: http://gomex.erma.noaa.gov/

[18] A. Clark, R. N. Swayze, I. Leifer, K. E. Livo, R. Kokaly, T. Hoefen, S. Lundeen, M. Eastwood, R. O. Green, N. Pearson, C. Sarture, I. McCubbin, D. Roberts, E. Bradley, D. Steele, T. Ryan, and R. Dominguez, "A method for quantitative mapping of thick oil spills using imaging spectroscopy," USGS Open File Rep. 2010-1167, Airborne Visible/Infrared Imaging Spectrometer (AVIRIS), 2010, pp. 2–7.

[19] L. Corucci, F. Nardelli, and M. Cococcioni, "Oil spill classification from multi-spectral satellite images: Exploring different machine learning techniques," *Proc. SPIE—Int. Soc. Opt. Eng.*, vol. 7825, DOI: 10.1117/12.864556782509.

[20] D. J. Lary and A. Walker, "High resolution identification of dust sources using machine learning and remote sensing data," presented at the Fall Amer. Geophys. Union Meeting, San Francisco, CA, 2011, AGU session A08-6.

ABOUT THE AUTHORS

Oleg Aulov received the B.S. degree in mathematics from the University of Central Missouri, Warrensburg, in 2004 and the M.S. degree in computer science with a concentration in computer security and information assurance from George Washington University, Washington, DC, in 2006. He is currently working toward the Ph.D. degree in the Department of Computer Science and Electrical Engineering, University of Maryland Baltimore County, Baltimore.

His topics of interest include social media mining, machine learning, trust establishment and management, computer security, information assurance, smart agents, and social engineering.

Milton Halem received the B.S. degree in mathematics from the City College of New York, New York, in 1952 and the Ph.D. degree in mathematics from the Courant Institute of Mathematical Sciences, New York University, New York, in 1968.

He has been a Research Professor in the Computer Science and Electrical Engineering Department, University of Maryland Baltimore County (UMBC), Baltimore, since 2002. In addition, he is the Director of the Multicore Computational Center and Principal Investigator of the National Science Foundation (NSF)-sponsored Center for Hybrid Multicore Productivity Research of the College of Engineering and Information Technology. Prior to joining UMBC, he served from 1999 to 2002 as Assistant Director for Information Sciences and Chief Information Officer at the NASA Goddard Space Flight Center and from 1984 to 1999 as Chief of the Earth and Space Data Computing. He was responsible for the management and conduct of one of the world's most powerful scientific data-intensive supercomputing complexes. His main teaching and research areas of current interest are service-oriented cloud computing, hybrid computational science, advanced information systems, and data-intensive computing.

Dr. Halem is nationally recognized for his research in simulation studies of space observing systems and for the development of 4-D data assimilation for weather and climate prediction. In 1999, he was awarded the honorary Doctorate degree from Dalhousie University, Halifax, NS, Canada, in recognition for his contributions to the field of computational science.