

Weather Radar Spatiotemporal Saliency: A First Look at an Information Theory–Based Human Attention Model Adapted to Reflectivity Images

DAVID SCHVARTZMAN AND SEBASTIÁN TORRES

Cooperative Institute for Mesoscale Meteorological Studies, University of Oklahoma, and NOAA/OAR/National Severe Storms Laboratory, Norman, Oklahoma

TIAN-YOU YU

School of Electrical and Computer Engineering, and Advanced Radar Research Center, and School of Meteorology, University of Oklahoma, Norman, Oklahoma

(Manuscript received 6 May 2016, in final form 6 October 2016)

ABSTRACT

Forecasters often monitor and analyze large amounts of data, especially during severe weather events, which can be overwhelming. Thus, it is important to effectively allocate their finite perceptual and cognitive resources for the most relevant information. This paper introduces a novel analysis tool that quantifies the amount of spatial and temporal information in time series of constant-elevation weather radar reflectivity images. The proposed Weather Radar Spatiotemporal Saliency (WR–STS) is based on the mathematical model of the human attention system (referred to as saliency) adapted to radar reflectivity images and makes use of information theory concepts. It is shown that WR–STS highlights spatially and temporally salient (attention attracting) regions in weather radar reflectivity images, which can be associated with meteorologically important regions. Its skill in highlighting current regions of interest is assessed by analyzing the WR–STS values within regions in which severe weather is likely to strike in the near future as defined by National Weather Service forecasters. The performance of WR–STS is demonstrated for a severe weather case and analyzed for a set of 10 diverse cases. Results support the hypothesis that WR–STS can identify regions with meteorologically important echoes and could assist in discerning fast-changing, highly structured weather echoes during complex severe weather scenarios, ultimately allowing forecasters to focus their attention and spend more time analyzing those regions.

1. Introduction

The network of Weather Surveillance Radar-1998 Doppler (WSR-88D) consists of 158 high-resolution, S-band, Doppler polarimetric weather radars operated by the U.S. National Weather Service (NWS) (Whiton et al. 1998), the U.S. Air Force, and the Federal Aviation Administration. The WSR-88D surveils the atmosphere by mechanically rotating a parabolic antenna using one of several predefined scanning strategies known as volume coverage patterns (VCP). VCPs for nonclear-air conditions take 4–6 min to complete, which defines the temporal resolution of the WSR-88D network. Forecasters of the NWS rely heavily on data from the WSR-88D to observe weather phenomena, and more

specifically to identify potentially severe weather (Andra et al. 2002). Storms may have remarkably different structures and can evolve in unpredictable ways. In particular, severe storms (e.g., supercell thunderstorms or hailstorms) can evolve in a matter of a few minutes (Heinselman et al. 2008) and potentially result in the loss of lives and property. Recent experiments strongly suggest that radar data with high temporal resolution could be beneficial in the warning decision process of NWS forecasters, leading to increased warning lead times for severe storms (Bowden et al. 2015; Heinselman et al. 2012). However, radar data with higher temporal resolution would increase the already large amount of data that forecasters have to analyze.

Forecasters monitor and process data from many different sources during severe weather situations, which can be overwhelming. For example, during a typical tornado outbreak situation, numerous convective supercell storms

Corresponding author e-mail: David Schwartzman, david.schwartzman@noaa.gov

can form with only a few of them producing tornadoes. This type of severe thunderstorm usually has distinctive radar signatures (hook echo, velocity couplets, etc.) that evolve quickly and attract the forecaster's attention. Attention is an instinctive biological mechanism that every living animal and human has; it enables us to selectively focus on the incoming stimuli and discard less interesting signals (Mancas et al. 2012). About 80% of the information we receive every day is from our visual system (Li and Gao 2014). The human retina can receive an equivalent of up to 10 billion bits of information per second (Raichle 2010); nevertheless, the cortex has only approximately 20 billion neurons (Shepherd 2003; Koch 2004). Therefore, the amount of information we receive significantly exceeds the amount of information we can actually store in our brains, and therefore we are routinely faced with information overload. This very same problem is exacerbated in weather forecasters who have to adopt a conceptual model of the atmosphere and analyze information from multiple sources (e.g., radar, satellite, surface observations, model output) simultaneously when making warning decisions. In addition to capacity limitations, our brains are constrained by their processing capabilities, and it becomes unattainable to simultaneously perform an efficient analysis of all the visual information (Li and Gao 2014). Even though our brains have these limitations, we are still able to accomplish highly dynamic tasks thanks to our attention system. Through the attention system, we can enormously reduce the amount of information flooding our brains and focus only on important sources of information. Attention can be defined as the allocation of cognitive resources on sources of relevant information. There are many types of attention, and they are usually classified by the way in which they make use of cognitive resources. According to Sohlberg and Mateer (1989), attention types can be classified as focused, sustained, selective, alternating, and divided. We are interested in the selective attention, which is defined as the quality to selectively maintain cognitive resources on a specific conspicuous object or region while ignoring all other competing stimuli (Li and Gao 2014). Focusing attention on relevant sources of information becomes crucial in overwhelming tasks, such as weather forecasting and warning.

A bio-inspired model of saliency (Itti et al. 1998, 2005; Tsotsos et al. 2005) has been applied to several fields (Frintrop et al. 2010; Itti and Koch 2001). For example, Li et al. (2011) used visual saliency to detect interesting regions in images and improve their compression by dynamically adjusting the resolution based on the degree of interest. Maddalena and Petrosino (2008) developed a saliency-based technique that separates foreground and background components for scenes of stationary cameras

used in video surveillance applications. In this work, saliency is used to highlight regions with high temporal and spatial information in weather radar reflectivity images. That is, the proposed Weather Radar Spatiotemporal Saliency (WR-STs) uses a computational model of the human's selective attention system based on information theory that seeks to resemble a forecaster's visual examination of weather radar reflectivity images. Such a mathematical model cannot capture the forecaster's complex conceptual model of the atmosphere; however, we postulate that WR-STs could help focus a forecaster's attention, since the structured and fast-evolving regions highlighted by WR-STs agree with regions of meteorological importance. Automatically identifying salient regions in weather radar images could be useful for several applications. For instance, WR-STs could assist human forecasters throughout the warning decision process as an additional nowcasting-like tool that highlights regions with higher saliency, especially in complex severe weather scenarios. Instead of performing manual observations of all the radar data collected in a volume scan, WR-STs could aid forecasters by eliminating the need to look at some elevation angles, thus increasing the time available for looking at important regions or data from other sources (e.g., weather stations and satellites). Furthermore, it could be applied to adaptive weather sensing (Reinoso-Rondinel et al. 2010; Torres et al. 2016), whereby regions with high saliency are updated more frequently than other less informative regions. These faster updates of quickly evolving, finely structured storm regions could aid in the interpretation of severe weather phenomena and could provide increased confidence in the radar data during the warning decision process.

The rest of the paper is organized as follows. Section 2 presents definitions and concepts associated with visual saliency and its activation function. In section 3, WR-STs is described mathematically, and its characteristics and implementation parameters are discussed. In section 4, data from a severe storm outbreak are processed using WR-STs, and the results are qualitatively analyzed. In section 5, the performance of the proposed metric is analyzed by correlating it with warning polygons issued by NWS forecasters during the development of the event. A quantitative analysis is carried out, and results are drawn from this discussion. Section 6 summarizes the conclusions of this work, reviews the limitations of the model, and provides recommendations for future work.

2. Visual saliency

The existence of an intuitive mechanism in the brain (and a visual map) that can depict conspicuous regions

in the field of view of the human visual system was originally proposed by Koch and Ullman (1987). In 1998, Itti and Koch developed the first computational implementation of a bottom-up, task-independent, selective visual attention model (Itti et al. 1998) that is the basis for most saliency models, including WR-STs. Their model has three main steps. First, three different visual features (color, intensity, and orientation) are computed from the input image. Second, a scale decomposition is applied to each feature map, and the outcome is a set of 12 feature maps (3 features, each with four scales). The third step is an across-scale combination followed by an activation (also called normalization) that highlights unique attributes from the feature maps. Last, a weighted linear combination of the three remaining activated-and-scale-combined features is made to arrive at the final spatial saliency map. This initial implementation of a selective attention model was very successful in providing better predictions of human eye fixations than simpler methods (e.g., chance or direct application of entropy), and inspired many research efforts that resulted in several saliency models. Different saliency models involve very different technical approaches, but they are all derived from the central concept of information *innovation* in some context (Riche et al. 2013). Popular saliency models include the Local and Global Saliency (LGS) by Itti et al. (1998), the Covariance Saliency (CovSal) by Erdem and Erdem (2013), the Graph-Based Visual Saliency (GBVS) by Harel et al. (2007), the Boolean Map-Based Saliency (BMS) by Zhang and Sclaroff (2013), the bottom-up algorithm for global rare feature detection (RARE2012) by Riche et al. (2013), and the Attention Based on Information Maximization (AIM) by Bruce and Tsotsos (2007). There are several technical differences among these methods, but the main difference is in their activation functions. As we shall discuss in more detail in the next section, the purpose of the activation function is to pick out unique attributes from the feature maps. LGS uses a deterministic function as the activation. CovSal uses a covariance-based activation technique to obtain the activated feature maps. GBVS constructs fully connected graphs over each feature map and assigns weights between nodes of the graph. Graphs are then treated as Markov chains to compute the activation maps, and maps are combined across features to obtain the final map. BMS computes saliency using the Boolean map theory of visual attention (Huang and Pashler 2007) with Bayesian-like activation functions. Finally, both RARE2012 and AIM use information theory type of activation functions (i.e., Shannon's entropy).

Saliency models can be divided into two broad categories based on their applications: bottom-up and top-down models. In bottom-up saliency models, attention is driven by salient stimuli without preassumptions (i.e., memory free), independent from prior knowledge. Bottom-up models use low-level features extracted from the image, such as intensity, contrast, and orientation. Once those features are generated, the models look for rare, contrasted, novel, more informative (less compressible) regions that maximize the amount of visual information. In other words, they all intend to find unusual characteristics in a given context: in the spatial, temporal, or both dimensions. On the other hand, saliency can also be guided by memory-dependent mechanisms with the use of prior or learned knowledge that is ingested into the model. These are called top-down saliency models, and they can dramatically increase the performance due to the prior knowledge incorporated in the model. It could be argued that the features of different weather phenomena in radar images may provide useful prior knowledge in top-down saliency models (e.g., it is more likely to observe severe thunderstorms in areas of high reflectivity). Similarly, task-specific features could be extracted to improve the detection of different types of salient areas (e.g., snow, freezing rain, and even convergence lines). However, the features observed by weather radar can widely vary depending on the radar acquisition parameters (e.g., spatial sampling, temporal resolution, variance of estimates) and the characteristics of the storms (e.g., intensity, orientation, distance from the radar). Thus, this proof-of-concept implementation adopts the simpler bottom-up saliency model; performance improvements based on top-down saliency models are left for future work.

As mentioned before, the application of visual saliency can be used to distribute the finite resources on the most relevant regions of an image. The spatial activation of the feature maps is a key step to accomplish this. Figure 1 (adapted from Itti et al. 1998) depicts this concept through an example in which the visual stimulus contains only a single salient feature: the red line oriented vertically among many horizontally oriented red lines. Each red line, regardless of its orientation, produces a maximum value in the intensity feature extraction. As a result, the intensity map consists of many peaks with the same value. In this case, it is expected that the activation of the intensity map will lead to a result indicating no information. On the contrary, the orientation map has a slightly higher peak in the region around the vertically oriented line. It is expected that the activation function can amplify this peak while suppressing the others. In other words,




FIG. 1. Notion of activation function adapted from Itti et al. (1998). (left) The visual stimulus used as an input for the model. After extracting the features, (top center) the intensity map has multiple equally high peaks, (top right) which after activation produce a flat map. In contrast, (bottom center) the orientation feature map shows a larger peak embedded in slightly smaller peaks, and (bottom right) activation enhances that peak while suppressing all others.

an ideal activation function would highlight only remarkable peaks from each feature map, picking out the most spatially informative regions. A conceptually similar activation can be performed in the time domain.

The computation of spatial and temporal saliency for weather radar reflectivity images is described in the next section.

3. WR-STTS

A simplified functional block diagram of the WR-STTS computation process is shown in Fig. 2. The blocks of this simplified diagram are described in detail next.

a. Multiscale decomposition and feature extraction

After interpolating the polar data collected by the radar into a Cartesian coordinate system (see Schwartzman 2015), we apply a multiscale decomposition and a spatial feature extraction as illustrated in step 1 of Fig. 3. Storm-scale and mesoscale weather phenomena can span a large range of spatial scales, with scales as small as 30 m for F0 tornadoes (Brooks 2004) to an extension of more than 100 km for very large thunderstorms or squall lines. The multiscale decomposition is accomplished

using the steerable pyramids method (Freeman and Adelson 1991; Simoncelli and Freeman 1995), which performs a linear multiscale, multiorientation image decomposition. In a nutshell, a directional wavelet decomposition of the Fourier transform of the image is used to obtain independent representations of scale and orientation that are translation and rotation invariant. The reader is referred to Mallat (2008) for more details about this technique.

The selection of the smallest scale is based on the average size of tornado widths reported on the ground for severe weather events with considerable damage (F2 or higher). Brooks (2004) reported that approximately 50% of F2 tornadoes have a mean width of 500 m. Thus, the smallest scale of this model is 500 m, and subsequent scales are obtained by doubling the size of the previous one. To avoid spatial aliasing, the resolution of the Cartesian radar grid is 250 m. The use of smaller scales might be problematic, since it could result in undesired echoes (clutter, biological scatterers, noise, etc.) being highlighted. It was determined (Schwartzman 2015) that using seven scales ($\gamma_1 = 500$ m, $\gamma_2 = 1$ km, $\gamma_3 = 2$ km, $\gamma_4 = 4$ km, $\gamma_5 = 8$ km, $\gamma_6 = 16$ km, and $\gamma_7 = 32$ km) is sufficient for identifying regions of interest in weather radar reflectivity images.




FIG. 2. Functional block diagram outlining the main steps for the computation of WR-STTS.

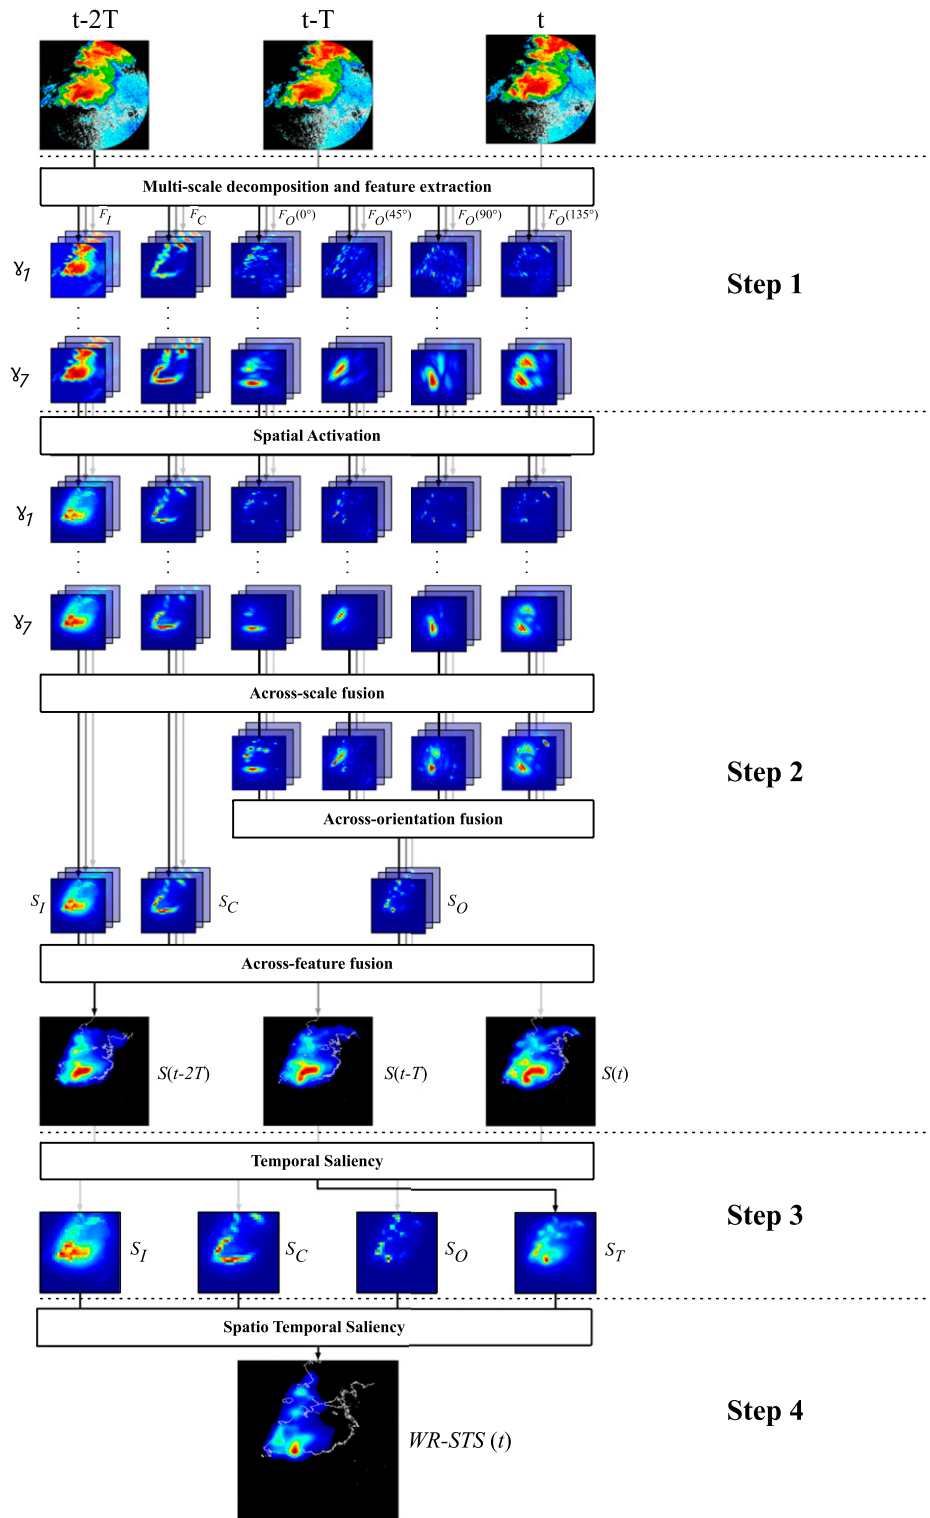


FIG. 3. Illustration of the steps for the computation of WR-STIS (reduced number of scales and orientation filters is used for simplicity). First, consecutive radar reflectivity images are scale decomposed, and their intensity, contrast, and orientation features are extracted. Then, a spatial activation function is applied on each scale, and single-feature maps are combined across all scales. Next, the feature maps are fused into spatial saliency maps, which are used to compute the temporal feature. Finally, the four feature maps are combined through a weighted average to produce the WR-STIS map.

This choice of scales will be supported by the results discussed in section 5.

Once the image is represented at multiple scales (as exemplified in Fig. 3 for two spatial scales), the feature extraction step is performed. Three types of spatial features are extracted at each scan time t , scale γ_l , and, if applicable, orientation θ_j ; these are denoted by $F_I(t, \gamma_l)$ for intensity, $F_C(t, \gamma_l)$ for contrast, and $F_O(t, \gamma_l, \theta_j)$ for orientation. To simplify the notation and without loss of generality, we assume a constant update time T such that t is an integer multiple of T . In this work, only the field of radar reflectivity is used, which can be considered as a grayscale image. The reflectivity is expressed in dBZ units because the logarithmic scale allows the model to highlight relatively weaker but important features that otherwise might be obscured. The intensity feature $F_I(t, \gamma_l)$ for grayscale images is simply the grayscale image matrix normalized in the range of $[0, 1]$, corresponding to the fixed range of possible reflectivity values, namely, $[-32, 94.5]$ dBZ. The contrast feature is computed by running a 2D window over every pixel of the image, taking a local neighborhood of pixels and computing the standard deviation of the pixel values in the neighborhood. For symmetry purposes, the neighborhood window used here is a disk, and edge effects are handled by replication. For the image with the highest resolution, a 10-pixel-diameter disk is used. For the next scale, the size of the window is reduced by one pixel in each dimension. The same is done subsequently to coarser scales to compute the contrast feature. The example in Fig. 3 shows only the first and last scales and four orientations; however, the WR-STs implementation uses seven scales and 16 orientations (Schvartzman 2015).

The orientation features can be derived from the reflectivity image at each scale by convolving it with Gabor filters oriented in specific directions. These Gabor filters can approximate the receptive field impulse response of orientation-selective neurons in the primary visual cortex (Daugman 1985); therefore, they are particularly suited for detecting features oriented in a particular direction that visually stand out from their surroundings. The convolution (denoted by \star) between the scale-decomposed reflectivity images and the Gabor filter can be computed as (Li and Gao 2014)

$$F_O(t, \gamma_l, \theta_j) = |F_I(t, \gamma_l) \star G_R(\theta_j)| + |F_I(t, \gamma_l) \star G_I(\theta_j)|, \quad (1)$$

where $\theta_j = j(180^\circ/16)$ is the j th orientation. A Gabor filter is the product of a sinusoid and a two-dimensional Gaussian function; the real and imaginary parts of the

complex Gabor filter (G_R and G_I , respectively) applied to the pixel (x, y) are defined as

$$G_R(\theta_j)(x, y) = \exp\left(-\frac{x^2 + y^2}{2}\right) \cos\left[\frac{2\pi}{\lambda}(x \cos\theta_j + y \sin\theta_j)\right] \quad (2)$$

and

$$G_I(\theta_j)(x, y) = \exp\left(-\frac{x^2 + y^2}{2}\right) \sin\left[\frac{2\pi}{\lambda}(x \cos\theta_j + y \sin\theta_j)\right], \quad (3)$$

where λ is the wavelength. The wavelengths used in WR-STs are $\lambda = 1, 2, 4, 8, 16, 32,$ and 64 , progressively doubling to match the choice of scales. In other words, because the pixel resolution is the same at each scale, the filter kernel must be adjusted to match the corresponding scale. A thorough description and characterization of these filters are provided in Feichtinger and Strohmer (1998).

b. Spatial activation and fusion

Following the scale decomposition and feature extraction, the spatial activation is carried out. As discussed previously, the choice of activation function is key in the performance of WR-STs, since its purpose is to pick only uniquely salient attributes from the feature maps. Because of the uncertainty typically present in weather phenomena, the use of deterministic functions would produce inconsistent results. Statistical methods use covariances as the main tool for activation and are usually very sensitive to contrast (large variances), while information theory activations produce more statistically robust results. The generality of the Shannon information metrics makes this class of functions suitable for weather radar images, since these can evolve into complex configurations that require statistically robust activations.

Shannon's entropy is defined in terms of probability distributions and has many properties that agree with the intuitive notion of what a measure of information should be. In fact, the entropy is considered to be the self-information of a signal. It can be shown that uniform probability functions have maximum entropy, since the outcome of a realization of the experiment is equiprobable for all the possible values of that random variable. On the other hand, narrower probability density functions are associated with less information, since the outcome of the corresponding experiment is more predictable.

The concept of entropy in a 2D image is exemplified in Fig. 4 using a radar reflectivity image. The region with




FIG. 4. Illustration of windowed computation of entropy for a weather radar reflectivity image (or intensity feature F_I). The entropy is a functional of the distribution and therefore depends only on the estimated probability density function, $p(F_I)$.

similar reflectivity values produces a relatively narrow probability density function (top-left box) and therefore can be thought of as having a more predictable spatial structure. On the other hand, in a region with large reflectivity variation, the probability density function (pdf) approximates a uniform distribution (bottom-left box), implying there is more uncertainty (or more information) in that region. For WR-STS, feature activation is carried out by computing the entropy of all the feature maps. That is, at each pixel of a feature map, the pdf is approximated with a 50-bin, unit-summation normalized histogram of the values inside a disk-shaped window (same as the one used for the contrast feature described previously). The entropy is computed as (Cover and Thomas 2006)

$$S = - \sum_f p(f) \log_2 p(f), \tag{4}$$

where $p(f)$ is the normalized histogram of the feature values inside the window.

The activation step produces one single-scale, single-feature saliency map for each feature map: $S_I(t, \gamma_I)$ for intensity, $S_C(t, \gamma_I)$ for contrast, and $S_O(t, \gamma_I, \theta_j)$ for orientation. These are fused across all scales to obtain single-feature saliency maps: $S_I(t)$ for intensity, $S_C(t)$ for contrast, and $S_O(t)$ for orientation (step 2 in Fig. 3).

To fuse the scales, first we linearly interpolate the images to the grid with the highest resolution. Then, for a given feature map, the pixel with maximum value is taken across all scales and the result is the single-feature spatial saliency map. Taking the maximum ensures that single-scale salient regions are captured when the multiscale images are aggregated. Single-orientation maps are also activated through this procedure and are then combined across orientations by averaging, since no particular orientation is favored. Finally, the spatial saliency map is obtained by a weighted average of the three single-feature spatial saliency maps as

$$S(t) = \frac{4S_I(t) + S_C(t) + 2S_O(t)}{7}. \tag{5}$$

The single-feature-map weights were fine-tuned in an ad hoc way with the assistance of a NWS forecaster. The forecaster was presented with three different weather events (a snowstorm, a squall line, and a tornadic supercell) and reviewed the reflectivity images from three consecutive scans for each event. Then, with a reasonable justification, the forecaster manually selected the most meteorologically important regions. Based on a qualitative interpretation of the relevance of each feature, we adjusted the weights until the WR-STS

maps roughly agreed with the forecaster's assessment. Whereas this simplistic approach may be subject to interpretation bias, we consider it sufficient for this initial proof-of-concept implementation. Through a more systematic and extensive process involving several forecasters and other weather events, these weights could be further refined. We leave this for future work.

c. Temporal activation

After obtaining the spatial saliency map, the next step is the temporal activation (step 3 in Fig. 3). In this work, mutual information (MI) is proposed as a means to characterize the temporal evolution of weather echoes in radar images. MI extends the notion of entropy and measures the information one random variable contains about another. The mutual information $M[S(t), S(t-T)]$ of current and previous spatial saliency maps $S(t)$ and $S(t-T)$ is computed as (Cover and Thomas 2006)

$$M[S(t), S(t-T)] = \sum_{s_t} \sum_{s_{t-T}} p(s_t, s_{t-T}) \log_2 \frac{p(s_t, s_{t-T})}{p(s_t)p(s_{t-T})}, \quad (6)$$

where $p(s_t)$ and $p(s_{t-T})$ are the marginal pdfs, and $p(s_t, s_{t-T})$ is the joint pdf. As before, the pdfs are

$$S_T(t) = 1 - \sum_{s_t} \sum_{s_{t-T}} \sum_{s_{t-2T}} p(s_t, s_{t-T}, s_{t-2T}) \log_2 \frac{p(s_t, s_{t-T})p(s_t, s_{t-2T})p(s_{t-T}, s_{t-2T})}{p(s_t, s_{t-T}, s_{t-2T})p(s_t)p(s_{t-T})p(s_{t-2T})}, \quad (7)$$

where the second term on the right-hand side of this equation is the MI of the current and two previous spatial saliency maps. As intended, this expression produces high values of temporal saliency for low values of MI (and vice versa) because the MI is low in regions with high temporal variations.

d. Spatiotemporal fusion

Although the temporal saliency map uses the spatial saliency maps obtained from the weighted single-feature spatial saliency maps, MI is a functional of the pdfs and thus independent of the weights. Therefore, as a means to fine-tune the performance of WR-STS by allowing independent weighting of the single-feature saliency maps, the spatiotemporal saliency map is computed as a weighted average of S_I , S_C , S_O , and S_T as (step 4 in Fig. 3)

$$\text{WR-STS}(t) = \frac{4S_I(t) + S_C(t) + 2S_O(t) + 2S_T(t)}{9}. \quad (8)$$

As mentioned before, these weights were experimentally selected to achieve good agreement between regions

TABLE 1. Summary of the model parameters used in WR-STS.

Parameter	Value/name
Interpolated grid size	1000 × 1000
Number of scales	7 (500 m, 1 km, 2 km, 4 km, 8 km, 16 km, 32 km)
Number of orientations	16 (0°, 11.25°, 22.5°, 33.75°, 45°, ..., 168.75°)
Spatial activation	Entropy
Temporal activation	Mutual information
Intensity map weight	4
Contrast map weight	1
Orientation map weight	2
Temporal map weight	2

approximated by normalized histograms of values in running windows positioned over corresponding pixels in the two saliency maps (i.e., the windows enclose the same geographical area in both maps). The expression $M[S(t), S(t-T)]$ is the reduction in the uncertainty of $S(t)$ due to the knowledge of $S(t-T)$. Note that if $M[S(t), S(t-T)] = 0$, then there is no common information between the saliency maps at different times; that is, each saliency map conveys unique information. To increase the robustness of the model, temporal saliency is computed as

highlighted by WR-STS and those of meteorological importance defined by a NWS forecaster. The map obtained in the last step is the WR-STS map at time t , which is designed to highlight regions with high spatial and temporal changes. The WR-STS model parameters discussed in this section are summarized in Table 1.

4. Application of WR-STS to a tornado outbreak in central Oklahoma

As reported by the NWS Weather Forecast Office, "a tornado outbreak occurred over parts of northern and central Oklahoma during the day on 24 May 2011, with violent tornadoes devastating several communities. By the end of the day, one EF-5, two EF-4, and two EF-3 tornadoes destroyed buildings, ripped up trees and power poles, and unfortunately, resulted in 11 deaths and 293 injuries" (NWS 2011). This weather event produced a total of 12 tornadoes in Oklahoma. In particular, an EF-5 tornado that developed and touched down in Canadian County became the strongest of them,

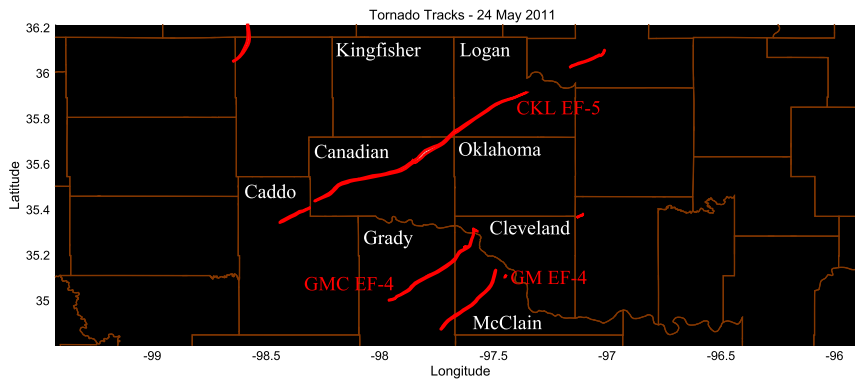


FIG. 5. Damage path left by 6 of the 12 tornadoes (the strongest ones) on 24 May 2011. County names mentioned in the text are shown for reference.

traveling through north of El Reno, west of Piedmont, and across the north edge of Guthrie. It was on the ground for almost 2 h (2050–2235 UTC) with a total pathlength of 101 km and a maximum width of 1.6 km. Nine people died as a result of this violent tornado and over 180 were injured. The damage path of all of these devastating tornadoes is shown in Fig. 5 (tracks of the less damaging tornadoes were not surveyed and are not displayed in the figure). The longest track is the one that corresponds to the Canadian–Kingfisher–Logan EF-5 (CKL EF-5) tornado. The two damage paths to the south of the CKL EF-5 tornado path correspond to the Grady–McClain–Cleveland EF-4 (GMC EF-4) and the Grady–McClain EF-4 (GM EF-4) tornadoes. These caused the death of one person and injured over 100 others.

By their nature, these damaging severe storms are intense (strong reflectivity cores) and can evolve quickly in time. It is expected that these types of convective storms exhibit more features than nonconvective storms, and that WR-STS will likely highlight structured, quickly-evolving regions. To confirm this, WR-STS is computed for the 0.5° -elevation reflectivity images, and severe weather warning polygons issued throughout the event are superimposed to assess their correlation. This makes sense because warning polygons indicate regions (of interest) where severe weather is present or is likely to strike in the near future. However, if the lead time of a warning polygon is less than two update times ($2T$), then WR-STS and warning polygons may not be completely independent. Still, even in situations where WR-STS and warning polygons are based on common reflectivity data, whereas WR-STS is based on three consecutive reflectivity images only, the determination of warning polygons involves multiple data sources (e.g., radar, satellite, weather stations) and, more importantly, the constantly evolving conceptual model of the forecaster at

the time of the warning (Waters 2007). Thus, good correlation between warning polygons and WR-STS can be used to infer the effectiveness with which WR-STS captures regions with meteorologically important echoes. Figure 6 shows reflectivity images spaced about 30 min apart (although WR-STS is computed on scans spaced approximately 5 min apart). The yellow superimposed polygons are severe thunderstorm warnings that were active at the time of the scan, and the red polygons are tornado warnings that were also active at the time. As a visual reference, white contour lines delimiting regions with 0.5 or higher WR-STS values are plotted on top of the reflectivity images. The corresponding WR-STS maps are presented in Fig. 7, with white 30-dBZ reflectivity contour lines plotted on top.

At 1932:07 UTC (Figs. 6a and 7a), many convective storm cells developed in western Oklahoma at about 160 km from the KTLX radar (Oklahoma City, near Twin Lakes). Two severe thunderstorm warnings were active at the time, although no tornado warnings had been issued at that point in time. One storm cell was not contained in a severe warning polygon at that time (farthest one northwest). Looking at the WR-STS map, it can be seen that both severe thunderstorm polygons correspond to medium-to-high WR-STS values, while the storm cell outside of the polygons corresponds to low WR-STS values. Notice that since the ground clutter is stationary (i.e., no temporal feature), its WR-STS is usually low (Schvartzman 2015). Later, at 1957:54 UTC, the storms continued to develop and two more severe thunderstorm warnings were issued (Figs. 6b and 7b) while the first two warnings remained active. Notice that even though the storms on the north have intense reflectivity cores, their WR-STS is relatively low compared to the southern storms, since the latter exhibit higher temporal features.




FIG. 6. Radar reflectivity scans of the tornadic storm on 24 May 2011 obtained from the KTLX radar. Active tornado warning polygons (red), severe thunderstorm polygons (yellow), and 0.5 or higher WR-STs contour lines (white) are overlaid on top of the reflectivity data: (a) 1932:07 (b) 1957:54 (c) 2027:43 (d) 2057:33 (e) 2127:20 (f) 2157:04 (g) 2226:48, and (h) 2256:33 UTC.

At 2027:43 UTC (Figs. 6c and 7c), two of the most recent severe thunderstorm warnings were still active, and a third one was issued as an update for the farthest storm located to the northwest. More importantly, three tornado warnings were issued and were active at the time of this scan. The storms for which these warnings were issued exhibit strong rotation features. WR-STs assigns very high values to the storm regions within the tornado warning polygons, with a slightly

higher value assigned to the southernmost storm with a maximum of 1. Notice that the WR-STs values within severe thunderstorm warnings are considerably lower than those in tornado warnings, corresponding to the notion that tornadic storms produce more structured and quickly evolving features.

Minutes before the data obtained by the scan shown in Fig. 6d, at approximately 2050 UTC, a tornado formed and became the CKL EF-5. Warnings were issued in

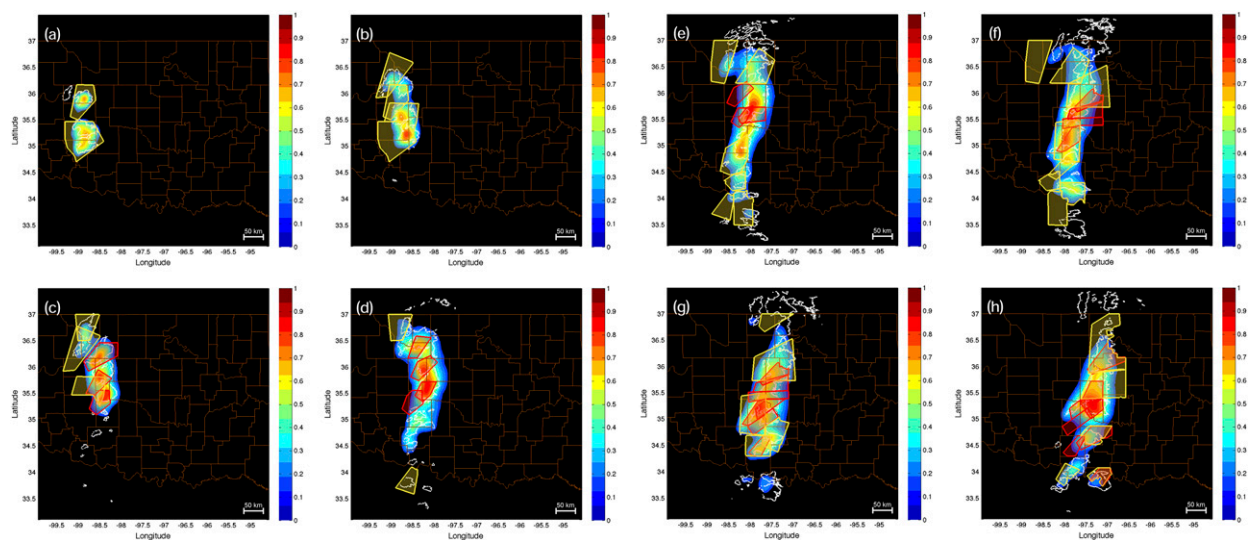


FIG. 7. Corresponding WR-STs maps for the reflectivity images in Fig. 6. Active tornado warning polygons (red), severe thunderstorm polygons (yellow), and 30-dBZ reflectivity contour lines (white) are overlaid on top of the WR-STs maps: (a) 1932:07, (b) 1957:54, (c) 2027:43, (d) 2057:33, (e) 2127:20, (f) 2157:04, (g) 2226:48, and (h) 2256:33 UTC.

west-central Oklahoma. Figures 6d and 7d shows that the previous three tornado warnings were still active and a fourth one was issued. WR-STS values are considerably higher in the last tornado warning polygon issued in Canadian County compared to the other three active warnings. In addition, a severe thunderstorm warning was issued for a storm cell to the southwest (outside of Oklahoma), and it is actually not deemed as informative by WR-STS. The reason for this is that the storm was very far from the radar and therefore its features are not apparent.

At 2127:20 UTC (Figs. 6e and 7e), a new tornado warning was issued for the CKL EF-5 tornado that was crossing Canadian County, and several new severe thunderstorm warnings were also issued. Once again, the WR-STS values assigned to weather echoes inside of the tornado warning polygons are noticeably higher than the values for severe thunderstorm warning polygons (which are also relatively high), even though most of the storm cells have comparable reflectivity values. At that moment, the CKL EF-5 tornado was still violent, and another funnel cloud started to form about 65 km to the south. Specifically, the storm that later became the GMC EF-4 tornado (west-central) was strengthening over the border between Caddo and Grady Counties, where WR-STS takes high values (around 0.75). Notice that the northernmost tornado warning polygon does not correspond with high WR-STS values, since this polygon was temporally obsolete at the time of the scan (it was issued over 30 min earlier; see Fig. 7d).

A tornado warning was issued at 2150 UTC for the GMC EF-4 tornado, as can be seen in Figs. 6f and 7f. At 2157:04 UTC, as the CKL EF-5 tornado was weakening over Logan County, WR-STS values also start decreasing, while they increase considerably in the region of the GMC EF-4 tornado that formed a little later, at approximately 2206 UTC. The analysis for Figs. 6g and 7g can be carried out in a similar way. Finally, at 2256:33 UTC (Figs. 6h and 7h), the CKL EF-5 tornado had dissipated completely and the GMC EF-4 tornado was starting to decay. WR-STS indicates a highly informative region over McClain and Cleveland Counties as the tornado crosses over. Because of the distance of the southern storm cells to the radar, many fine features are lost, and WR-STS computes only relatively low values in that region. This is perhaps one of the main limitations of WR-STS, reflecting the inherent limitation of weather radars to resolve small-scale features as storms get farther away.

5. Performance analysis and validation of WR-STS

With the goal of evaluating the statistical performance of WR-STS, we propose the use of archived severe thunderstorm and tornado warning polygons issued by

NWS forecasters. These polygons were determined by forecasters for regions in which hazardous weather would likely take place in the near future. Nevertheless, we stress that WR-STS is not designed to anticipate the occurrence of severe weather but to assist in highlighting regions where potentially hazardous weather is currently evolving.

High-reflectivity regions can be intuitively associated with regions of meteorological importance, since severe weather is often characterized by strong reflectivity cores. In fact, WR-STS uses reflectivity (intensity) as one of its spatial features. However, even though the intensity feature is an important spatial component in WR-STS, the other spatial and temporal components add significant value to its skill in highlighting important regions. In other words, regions of interest determined by WR-STS are not just high-reflectivity regions. To show this, the performance analysis is carried out on both WR-STS and radar reflectivities.

For the case study presented in section 4, radar reflectivity and WR-STS values inside and outside the areas determined by severe thunderstorm and tornado warning polygons are recorded for every scan throughout the time series, and normalized histograms (to an area of 1) are computed for both reflectivity and WR-STS (Fig. 8). An examination of these histograms reveals that it may be easier to distinguish regions with meteorologically important echoes using WR-STS, since the normalized histograms of values inside and outside warning polygons are more separated. To quantify the statistical distance between these normalized histograms, we use the *total variation distance*. Whereas this measure is related to the relative entropy (also known as the Kullback–Leibler divergence), it has the important property of being a distance metric on the space of probability distributions. It is defined as (Sriperumbudur et al. 2009)

$$\delta(I, O) = \frac{1}{2} \sum_{x \in \Omega} |P_I(x) - P_O(x)|, \tag{9}$$

where I and O are random variables over the same domain Ω , corresponding to the values inside and outside the warning polygons, respectively (referred to as “in” and “out,” respectively). These random variables have probability mass functions given by $P_I(x)$ and $P_O(x)$, respectively, where $x \in \Omega$. Notice that the term inside the summation on the right-hand side represents the absolute difference between probabilities of each random variable, resulting in the same outcome and it is bounded between 0 and 1. For instance, if I and O are statistically very similar, then their densities are almost overlapped and $\delta(I, O)$ will approach 0. On the other hand, if I and O are statistically very different, there will be little to no overlap between their densities, and




FIG. 8. Selective ability of WR-STs: (a) normalized histograms (unit area) of reflectivity inside and outside the warning polygons. (b) As in (a), but for WR-STs; (c) ROC curves of reflectivity and WR-STs; and (d) time-composite WR-STs.

$\delta(I, O)$ will approach 1. In the context of classification, the larger the total variation distance between classes, the better the separation between them. Thus, by comparing the total variation distance of WR-STs and the reflectivity inside and outside the polygons, we can assess their relative skills at highlighting meteorologically important regions. To this end, we proceed to determine the total variation, using normalized histograms of reflectivity and WR-STs.

We compute normalized (unit area) histograms of reflectivity and WR-STs to approximate their probability mass functions inside and outside the warning polygons. Denoting the normalized histograms of reflectivity as in and out by \hat{P}_{I_s} and \hat{P}_{O_s} , respectively, we compute $\delta(I_s, O_s) = 0.407$. This means that 40.7% of the area below the curves in Fig. 8a is not shared, while 59.3% is. Similarly, denoting the normalized histograms of WR-STs by \hat{P}_{I_s} and \hat{P}_{O_s} , we compute $\delta(I_s, O_s) = 0.709$. This indicates that 70.9% of the area below the curves in Fig. 8b is not shared, while 29.1% is. There is a substantial difference between the total variation distance of reflectivity and WR-STs, which confirms that the additional features incorporated into WR-STs make it more skilled than reflectivity only. These normalized histograms are used to compute receiver operating characteristics (ROC)

curves for reflectivity and WR-STs, which are shown in Fig. 8c, and to allow for a direct comparison of the relationship between the true and false positives of each technique. As the results show, WR-STs always attains a higher rate of true positives at a lower rate of false positives. For example, for a true positive rate of 90%, WR-STs has a false positive rate of approximately 20%, whereas reflectivity has a false positive rate of 65%. In addition to the normalized histograms and the ROC curves, a time-composite WR-STs map with superimposed warning polygons is shown in Fig. 8d. The map is computed by taking the maximum WR-STs throughout the time series for each resolution volume. Because WR-STs was designed to highlight regions with salient intensity, contrast, orientation, and temporal features, and tornadoes produce these kinds of features in radar reflectivity images, the highlighted path shown in Fig. 8d is, as expected, well correlated with tornado tracks.

It can be inferred from this case study that WR-STs could be used to automatically highlight regions of interest that contain meteorologically important echoes. Whereas the results presented above correspond just to a single severe weather event, they motivate the study of a larger number of cases. Significant severe weather

TABLE 2. WSR-88D reflectivity data from severe weather events used to evaluate the performance of WR-STs. Radars are as follows: KTLX near Twin Lakes, OK; KFWS near Fort Worth, TX; KHTX near Huntsville, AL; KOAX near Omaha, NE; KLSX near St. Louis, MO; KILX near Lincoln, IL; KLOT near Chicago, IL; and KEAX near Kansas City, MO.

	Radar	Date	Period (UTC)	Storm type
1	KTLX	20 May 2013	1904–2158	Isolated supercells
2	KFWS	3 Apr 2014	2115–0211	Isolated supercells
3	KLSX	7 Apr 2015	1203–1601	Isolated supercells
4	KLOT	22 Jun 2015	0003–0303	Isolated supercells
5	KHTX	28 Apr 2014	2003–0100	Squall line with multicells
6	KTLX	25 Mar 2015	2202–0212	Squall line with multicells
7	KILX	7 Jun 2015	2029–0101	Squall line with multicells
8	KEAX	6 Jul 2015	1901–2303	Squall line with multicells
9	KTLX	24 May 2011	1902–2356	Tornadic supercells
10	KOAX	3 Jun 2014	2000–2301	Tornadic supercells

events that occurred in the last few years in the Midwest United States were chosen for this study, where cases after 1 October 2007 were selected because storm-based polygon warnings became operational after that date.

Table 2 presents the basic characteristics of the selected severe weather events. For each event, a time series of reflectivity data from the lowest elevation scan (0.5°) was obtained from the National Climatic Data Center¹ (now known as NCEI; radar reflectivity) and the corresponding warning polygons from the NWS severe weather archives.² The second column in Table 2 indicates the name of the WSR-88D radar site. Dates and periods of the observations are specified in the following two columns. The last column specifies the overall type of severe weather event. It should be noted that the spatial and temporal characteristics of these severe storms are very different. More specifically, there are four isolated supercell cases (1–4), four squall lines with embedded storm cell cases (5–8), and two tornado outbreak cases (9 and 10).

Table 3 provides detailed information about the warnings issued throughout the period, as well as the performance obtained for reflectivity and WR-STs. *Polygon count* represents the number of thunderstorm warning polygons issued in the period, which were used to compute the performance metrics. *Warning avg size* (km^2) and *Warning avg duration* (min) express the average surface and average time span of the warnings polygons, respectively. Columns five and six provide an intuitive idea of the intensity of these storms through the average maximum reflectivity (dBZ) and average reflectivity core size (km^2), respectively, of the storms encompassed by the warning polygons. Both are taken at an elevation of 0.5° , where a reflectivity threshold of 35 dBZ is used to define the cores (Johnson et al. 1998). Notice that the tornado

outbreaks have the largest number of warnings issued. Moreover, they also have the largest average maximum reflectivity and average reflectivity core sizes. The number of warnings issued for squall lines and isolated supercells is variable, but the squall-line cases tend to have larger reflectivity cores and lower average reflectivity values than the isolated supercells due to their large spatial extent. Complex severe weather systems like these are of particular interest in this work since they help reveal the skill of WR-STs.

The total variation distance of the normalized histograms is given in columns seven and eight of Table 3. The results show that the total variation distance of WR-STs exceeds that of reflectivity for all the cases. Notice that the difference is always larger than 0.2, which shows that in and out WR-STs histograms are separated by at least 20% more than those corresponding to reflectivity. The mean total variation distance of reflectivity is 39.2%, while the mean for WR-STs is 71.2%. This is strong corroboration to the case study findings that WR-STs is particularly skilled in highlighting regions of meteorological importance characterized by warning polygons. In addition, notice that the lowest total variation distance for WR-STs is 0.5695 (case 5), which corresponds to a tornado outbreak case. This relatively low performance is due to the higher complexity of the storms in this case and the large number of warnings issued in the time period.

As mentioned before, warning polygons are regions determined by NWS forecasters in which severe weather is likely to strike. In this paper, warning polygons were used as ground truth to analyze the performance of WR-STs. However, the use of warning polygons as a validation method has some limitations. First, since the polygons are predicting the occurrence of severe weather, they do not always coincide with the region where severe weather actually takes place. In other words, in a few complex cases, the polygons may be incorrectly placed. Second,

¹ <http://ncdc.noaa.gov/>.

² <https://nwschat.weather.gov/lst/>.

TABLE 3. Warning polygon characteristics and performance metrics for the 10 severe weather cases in Table 2. Polygon count is the number of warning polygons active through the time period, avg. maximum reflectivity (dBZ) is the average maximum reflectivities recorded through the time period, and average core size (km²) is for the reflectivities above 35 dBZ in the 0.5° cut only. $\delta[I(z), O(z)]$ and $\delta[I(s), O(s)]$ are the total variation distances of the normalized histograms for reflectivity and WR-STs, respectively.

	Polygon count	Warning avg size (km ²)	Warning avg duration (min)	Avg max reflectivity (dBZ)	Avg core size (km ²)	$\delta[I(z), O(z)]$	$\delta[I(s), O(s)]$
1	12	1327.0	39.2	49.2	855.9	0.3527	0.5764
2	16	2227.2	32.5	61.1	1031.9	0.3327	0.7784
3	7	1706.7	43.2	60.1	730.2	0.5065	0.8231
4	6	1969.3	38.6	56.4	434.5	0.2384	0.6133
5	10	861.8	29.0	30.4	560.5	0.5485	0.7593
6	22	1973.1	38.5	55.9	1614.1	0.3540	0.7992
7	24	1475.2	35.9	52.9	2061.5	0.3987	0.6378
8	6	2277.6	43.6	37.8	895.1	0.4769	0.8584
9	27	2453.1	29.8	64.5	2402.4	0.4068	0.7086
10	31	2229.0	35.6	72.3	2949.7	0.3044	0.5695

because severe weather may develop quickly and because the temporal resolution of the precipitation VCPs is about 5 min, forecasters may miss the event and therefore no polygon is issued in time. According to Barnes et al. (2007), the national false alarm ratio (FAR) for tornado warnings in 2003 was of 0.76, meaning that only one out of every four tornadoes warnings were verified.

Despite these limitations, it can be inferred from this analysis that high WR-STs values are in good agreement with regions of meteorological importance as defined by NWS forecasters. However, WR-STs does not have the ability to anticipate the occurrence of severe weather (as do forecasters), but it could assist them in complex weather scenarios. A computer-assisted human decision-making process could reduce the time spent manually looking at radar data from each elevation angle and increase the time available for looking at the most important storm regions. In turn, having more time to examine those storms could provide more confidence and aid in the warning decision process.

6. Conclusions

This paper explored the first application of a bio-inspired model of attention (referred to as saliency) to weather radar reflectivity images. Saliency models have been used in many other fields to model the human attention system to better allocate the limited resources on relevant information. The proposed model—Weather Radar Spatiotemporal Saliency (WR-STs)—accounts for spatial and temporal features present in radar reflectivity images. The spatial features include the intensity, contrast, and orientation at a number of different spatial scales. Aside from displaying regions of high information content, spatial features highlight regions with high meteorological information. Temporal features are obtained

by processing time series of spatial-saliency maps using mutual information. In general, we postulate that WR-STs could aid forecasters in focusing their attention by spending more time analyzing regions of meteorological importance, herein defined as weather echoes confined in severe thunderstorm and tornado warning polygons issued by NWS forecasters. It should be noted that, depending on the application, regions of meteorological importance could be defined in different ways (e.g., snowstorms or drylines); we leave this extension of WR-STs for future work.

It was shown through a case study that convective storms, which display distinctive spatial and temporal features, lead to higher WR-STs values. The analysis of a convective severe weather environment in which multiple tornado-producing thunderstorms were present revealed consistency between high WR-STs values and severe thunderstorm and tornado warning polygons issued by NWS forecasters. It was shown that WR-STs generally assigns medium-to-high values to regions in which severe thunderstorm warnings are active. These results were corroborated by analyzing 10 diverse cases using the same methodology. That is, based on a statistical-distance metric, WR-STs was shown to be significantly better than reflectivity at highlighting regions of meteorological importance.

Whereas WR-STs appears to be a promising tool to analyze weather radar images, it is still in its infancy, and this initial proof-of-concept implementation has a few limitations. First, image artifacts, such as beam blockage and anomalous propagation clutter, can cause WR-STs to highlight regions not necessarily of meteorological importance. Second, a relatively high temporal resolution and good azimuthal sampling are needed by this model to better focus on regions of interest. Perhaps one of the most important limitations of WR-STs is

related to its diminishing capability as the range from the radar increases. As the range from the radar increases, the radar resolution volumes become larger and many spatial features of severe storms are inherently obscured. This results in a significant reduction of WR-STs values and could lead to missing meteorologically important regions that are far from the radar.

Even though preliminary results are promising, there are a number of research questions that need to be explored before WR-STs could be considered for operational use. First, the proof-of-concept WR-STs uses images from only the lowest elevation scan. Using volumetric data could provide relevant information about higher levels of the atmosphere, and they could aid WR-STs in more accurately narrowing down regions where severe weather phenomena are developing. For example, mesocyclones usually develop at midlevels (Stumpf et al. 1998) and their initiation would generally not be seen at the lower elevations in radar images. Volumetric saliency is an active, ongoing research field and it is in its developing stages (Shen et al. 2016). Second, the model presented in this work uses only radar reflectivity images. This can largely limit the extent to which WR-STs is able to identify features in the radar data, since several important salient features may be present in other radar fields (e.g., mesocyclone signatures in radial velocity images). Incorporating the Doppler moments (radial velocity and spectrum width) and the polarimetric variables could provide the model with more information, making WR-STs more robust. Last, a study involving forecasters using eye-tracking tools (Bowden et al. 2016) could be used to validate the performance of WR-STs more systematically and accurately.

Although the number of cases analyzed is limited, we can conclude that WR-STs has the potential to consistently highlight regions of meteorological interest. In particular, WR-STs could aid in discerning meteorologically important weather echoes during complex severe weather situations, assisting human forecasters in the warning decision process by reducing the time spent examining all available radar data, and increasing the time available for analyzing the most important storm regions.

Acknowledgments. The authors thank Donald Burgess and Katie Bowden for providing comments that improved the manuscript. Funding was provided by NOAA/Office of Oceanic and Atmospheric Research under NOAA–University of Oklahoma Cooperative Agreement NA11OAR4320072, U.S. Department of Commerce.

REFERENCES

- Andra, D. L., E. M. Quetone, and W. F. Bunting, 2002: Warning decision making: The relative roles of conceptual models, technology, strategy, and forecaster expertise on 3 May 1999. *Wea. Forecasting*, **17**, 559–566, doi:10.1175/1520-0434(2002)017<0559:WDMTRR>2.0.CO;2.
- Barnes, L. R., E. C. Grunfest, M. H. Hayden, D. M. Schultz, and C. Benight, 2007: False alarms and close calls: A conceptual model of warning accuracy. *Wea. Forecasting*, **22**, 1140–1147, doi:10.1175/WAF1031.1.
- Bowden, K. A., P. L. Heinselman, D. M. Kingfield, and R. P. Thomas, 2015: Impacts of phased-array radar data on forecaster performance during severe hail and wind events. *Wea. Forecasting*, **30**, 389–404, doi:10.1175/WAF-D-14-00101.1.
- , —, and Z. Kang, 2016: Exploring applications of eye-tracking in operational meteorology research. *Bull. Amer. Meteor. Soc.*, doi:10.1175/BAMS-D-15-00148.1, in press.
- Brooks, H. E., 2004: On the relationship of tornado path length and width to intensity. *Wea. Forecasting*, **19**, 310–319, doi:10.1175/1520-0434(2004)019<0310:OTROTP>2.0.CO;2.
- Bruce, N., and J. Tsotsos, 2007: Attention based on information maximization. *J. Vision*, **7**, 950, doi:10.1167/7.9.950.
- Cover, T. M., and J. A. Thomas, 2006: *Elements of Information Theory*. 2nd ed. John Wiley & Sons, 792 pp.
- Daugman, J. G., 1985: Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional spatial cortical filters. *J. Opt. Soc. Amer.*, **2A**, 1160–1169, doi:10.1364/JOSAA.2.001160.
- Erdem, E., and A. Erdem, 2013: Visual saliency estimation by nonlinearly integrating features using region covariances. *J. Vision*, **13**, 11, doi:10.1167/13.4.11.
- Feichtinger, H. G., and T. Strohmer, Eds., 1998: *Gabor Analysis and Algorithms: Theory and Applications*. Applied and Numerical Harmonic Analysis, Birkhäuser Basel, 496 pp., doi:10.1007/978-1-4612-2016-9.
- Freeman, W. T., and E. H. Adelson, 1991: The design and use of steerable filters. *IEEE Trans. Pattern Anal. Mach. Intell.*, **13**, 891–906, doi:10.1109/34.93808.
- Frintrop, S., E. Rome, and H. I. Christensen, 2010: Computational attention systems and their cognitive foundations: A survey. *ACM Trans. Appl. Percept.*, **7**, 6, doi:10.1145/1658349.1658355.
- Harel, J., C. Koch, and P. Perona, 2007: Graph-based visual saliency. *Advances in Neural Information Processing Systems 19: Proceedings of the 2006 Conference*, B. Schölkopf, J. Platt, and T. Hofmann, Eds., Neural Information Processing Series, MIT Press, 545–552.
- Heinselman, P. L., D. L. Prieznitz, K. L. Manross, T. M. Smith, and R. W. Adams, 2008: Rapid sampling of severe storms by the National Weather Radar Testbed Phased Array Radar. *Wea. Forecasting*, **23**, 808–824, doi:10.1175/2008WAF2007071.1.
- , D. S. LaDue, and H. Lazrus, 2012: Exploring impacts of rapid-scan radar data on NWS warning decisions. *Wea. Forecasting*, **27**, 1031–1044, doi:10.1175/WAF-D-11-00145.1.
- Huang, L., and H. Pashler, 2007: A Boolean map theory of visual attention. *Psychol. Rev.*, **114**, 599, doi:10.1037/0033-295X.114.3.599.
- Itti, L., and C. Koch, 2001: Computational modelling of visual attention. *Nat. Rev. Neurosci.*, **2**, 194–203, doi:10.1038/35058500.
- , —, and E. Niebur, 1998: A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, **20**, 1254–1259, doi:10.1109/34.730558.

- , G. Rees, and J. Tsotsos, Eds., 2005: Models of bottom-up attention and saliency. *Neurobiology of Attention*, Academic Press, 576–582.
- Johnson, J. T., P. L. MacKeen, A. Witt, E. D. Mitchell, G. J. Stumpf, M. D. Eilts, and K. W. Thomas, 1998: The Storm Cell Identification and Tracking algorithm: An enhanced WSR-88D algorithm. *Wea. Forecasting*, **13**, 263–276, doi:10.1175/1520-0434(1998)013<0263:TSCIAT>2.0.CO;2.
- Koch, C., 2004: *Biophysics of Computation: Information Processing in Single Neurons*. Computational Neuroscience Series, Oxford University Press, 588 pp.
- , and S. Ullman, 1987: Shifts in selective visual attention: Towards the underlying neural circuitry. *Matters of Intelligence: Conceptual Structures in Cognitive Neuroscience*, L. M. Vaina, Ed., Synthese Library, Vol. 188, Springer, 115–141.
- Li, J., and W. Gao, Eds., 2014: *Visual Saliency Computation: A Machine Learning Perspective*. Lecture Notes in Computer Science, Vol. 8408, Springer International Publishing, 240 pp., doi:10.1007/978-3-319-05642-5.
- Li, Z., S. Qin, and L. Itti, 2011: Visual attention guided bit allocation in video compression. *Image Vision Comput.*, **29**, 1–14, doi:10.1016/j.imavis.2010.07.001.
- Maddalena, L., and A. Petrosino, 2008: A self-organizing approach to background subtraction for visual surveillance applications. *IEEE Trans. Image Process.*, **17**, 1168–1177, doi:10.1109/TIP.2008.924285.
- Mallat, S., 2008: *A Wavelet Tour of Signal Processing: The Sparse Way*. 3rd ed. Academic Press, 832 pp.
- Mancas, M., D. De Beul, N. Riche, and X. Siebert, 2012: Human attention modelization and data reduction. *Video Compression*, A. Punchihewa, Ed., InTech, 103–128, doi:10.5772/34942.
- NWS, 2011: The May 24, 2011 tornado outbreak in Oklahoma. [Available online at <https://www.weather.gov/oun/events-20110524>.]
- Raichle, M. E., 2010: The brain's dark energy. *Sci. Amer.*, **302**, 44–49, doi:10.1038/scientificamerican0310-44.
- Reinoso-Rondinel, R., T.-Y. Yu, and S. Torres, 2010: Multifunction phased-array radar: Time balance scheduler for adaptive weather sensing. *J. Atmos. Oceanic Technol.*, **27**, 1854–1867, doi:10.1175/2010JTECHA1420.1.
- Riche, N., M. Mancas, M. Duvinage, M. Mibulumukini, B. Gosselin, and T. Dutoit, 2013: RARE2012: A multi-scale rarity-based saliency detection with its comparative statistical analysis. *Signal Process.: Image Commun.*, **28**, 642–658, doi:10.1016/j.image.2013.03.009.
- Schvartzman, D., 2015: Weather radar spatio-temporal saliency. M.S. thesis, University of Oklahoma, 152 pp.
- Shen, E., Y. Wang, and S. Li, 2016: Spatiotemporal volume saliency. *J. Visualization*, **19**, 157, doi:10.1007/s12650-015-0293-y.
- Shepherd, G., Ed., 2003: *The Synaptic Organization of the Brain*. 5th ed. Oxford University Press, 736 pp.
- Simoncelli, E. P., and W. T. Freeman, 1995: The steerable pyramid: A flexible architecture for multi-scale derivative computation. *Proceedings: International Conference on Imaging Processing*, Vol. 3, IEEE, 444–447, doi:10.1109/ICIP.1995.537667.
- Sohlberg, M. M., and C. A. Mateer, 1989: *Introduction to Cognitive Rehabilitation: Theory and Practice*. Guilford Press, 414 pp.
- Sriperumbudur, B. K., K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. G. Lanckriet, 2009: On integral probability metrics, ϕ -divergences and binary classification. arXiv: 0901.2698, 18 pp. [Available online at <https://arxiv.org/pdf/0901.2698v4.pdf>.]
- Stumpf, G. J., A. Witt, E. D. Mitchell, P. L. Spencer, J. Johnson, M. D. Eilts, K. W. Thomas, and D. W. Burgess, 1998: The National Severe Storms Laboratory mesocyclone detection algorithm for the WSR-88D. *Wea. Forecasting*, **13**, 304–326, doi:10.1175/1520-0434(1998)013<0304:TNSSLM>2.0.CO;2.
- Torres, S. M., and Coauthors, 2016: Adaptive-weather-surveillance and multifunction capabilities of the National Weather Radar Testbed phased array radar. *Proc. IEEE*, **104**, 660–672, doi:10.1109/JPROC.2015.2484288.
- Tsotsos, J. K., L. Itti, and G. Rees, 2005: A brief and selective history of attention. *Neurobiology of Attention*, L. Itti, G. Rees, and J. K. Tsotsos, Eds., Academic Press, xxiii–xxxii, doi:10.1016/B978-012375731-9/S0003-3.
- Waters, K. R., 2007: Verification of National Weather Service warnings using geographic information systems. Preprints, *23rd Conf. on IIPS*, San Antonio, TX, Amer. Meteor. Soc., 4B.1. [Available online at https://ams.confex.com/ams/87ANNUAL/techprogram/paper_116773.htm.]
- Whiton, R. C., P. L. Smith, S. G. Bigler, K. E. Wilk, and A. C. Harbuck, 1998: History of operational use of weather radar by U.S. weather services. Part II: Development of operational Doppler weather radars. *Wea. Forecasting*, **13**, 244–252, doi:10.1175/1520-0434(1998)013<0244:HOOUOW>2.0.CO;2.
- Zhang, J., and S. Sclaroff, 2013: Saliency detection: A Boolean map approach. *2013 IEEE International Conference on Computer Vision (ICCV 2013)*, IEEE, 153–160, doi:10.1109/ICCV.2013.26.