

# Design and Evaluation of Trust–Eliciting Cues in Drug–Drug Interaction Alerts

DEBALEENA CHATTOPADHYAY<sup>1</sup>, NITYA VERMA<sup>2</sup>, JON DUKE<sup>3</sup>  
AND DAVIDE BOLCHINI<sup>2,\*</sup>

<sup>1</sup>*Department of Computer Science, University of Illinois at Chicago, 851 S. Morgan (M/C 152), Chicago, IL 60607, USA*

<sup>2</sup>*Department of Human-Centered Computing, Indiana University School of Informatics and Computing, 535 W Michigan St, Indianapolis, IN 46202, USA*

<sup>3</sup>*College of Computing, Georgia Institute of Technology, 75 5th St. Atlanta, GA 30308, USA*

*\*Corresponding author: dbolchin@iupui.edu*

---

**Drug–drug interaction (DDI) alerts safeguard patient safety during medication prescribing, but are often ignored by physicians. Despite attempts to improve the usability of such alerts, physicians still mistrust the relevance of simplistic computerized warnings to support complex medical decisions. By building on prior fieldwork, this paper evaluates novel designs of trust–eliciting cues in DDI alerts. A sequential mixed-method study with 70 physicians examined what trust cues improve compliance, promote reflection, and trigger appropriate actions. In a survey, 52 physicians rated the likelihood of compliance and usefulness of redesigned alerts. Based on these findings, alerts were assessed in a scenario-based simulation with 18 physicians prescribing medications in 6 patient scenarios. Our results show that alerts embodying expert endorsement, awareness of prior actions, and peer advice were less likely to be overridden than current alerts, and promoted reflection, monitoring, or order modifications—thus building towards greater attention to patient safety.**

## RESEARCH HIGHLIGHTS

- In a mixed-method sequential study, we designed and evaluated 11 types of DDI alerts embodying different trust–eliciting cues.
- Trust cues included expert endorsements (Chief of Quality, Institute of Medicine, Literature, Department Head, Institutional History, Endorsed-empathy), history, collaboration, empathy, and transparency.
- The survey indicates physicians were more likely to override the *control* alert than *agency-laden*, *endorsed*, and *collaborative* alert.
- The scenario-based lab study with physicians indicated alerts embodying expert endorsement with data, awareness of prior actions, and peer advice were less likely to be overridden than current DDI alerts.
- We provide holistic metrics that go beyond the binary cancel/continue compliance measure.

*Keywords: healthcare information systems; health informatics; clinical alerts; trust in socio-technical systems; interaction design; decision-making*

*Editorial Board Member: Dr. Fabio Paternò*

*Received 18 April 2017; Revised 2 October 2017; Editorial Decision 12 December 2017; Accepted 18 December 2017*

---

## 1. INTRODUCTION

In this age of computerized physician order entry (CPOE) systems, computerized clinical alerts pervade a physician's work of medication prescribing. Especially, drug–drug interaction (DDI) alerts warn physicians about the potential adverse effects of a certain drug combination on a patient. Upon receiving an alert, physicians decide to either cancel the drug order, thus complying with the warning, or continue.

Recent reports indicate alert compliance is at an all-time low; the rates of alert override rose from about 88% in 2002 (Payne *et al.*, 2002) to 95% in 2014 (Bryant, Fletcher and Payne, 2014). In part, the low rate of compliance may be attributable to alert fatigue, with 10–36% of all medication orders leading to some type of warning (van der Sijs *et al.*, 2008). But alert fatigue alone is insufficient to account for the near total disregard of computerized recommendations. Another important reason is the physician's lack of trust in electronic guidance, as noted in prior studies (Alexander, 2006; Isaac *et al.*, 2009). As a practice, medicine lacks general rules that can be unambiguously applied to every case at hand (Hunter, 1996). This leads physicians to remain firmly committed to personal or trusted experiences over textbook rules or the medical literature (Alexander, 2006; Greenhalgh, 1999). But alert overrides carry significant consequences, with patients experiencing serious adverse events as a result (Duke, Li and Dexter, 2013).

Prior research investigated ways to combat such alarming alert override rates. For example, by improving the drug knowledge bases to increase an alert's positive predictive value (van der Sijs *et al.*, 2008) or drawing on human-factor principles (Enkin and Jadad, 1998; Russ *et al.*, 2014) to improve an alert's layout, timeliness, and overall usability. But these efforts had limited success in increasing alert compliance, prompting recent work in two other directions.

First, observational studies have examined how physicians structure their daily decisions on medication prescribing around mutual trust, through peer discussions and by attending advice from their colleagues (Chattopadhyay *et al.*, 2015). The key elements of trusted peer advice, such as expert endorsements, opportunities to talk to colleagues about a decision at hand, awareness of prescribing history, or references to relevant literature, were codified into the trusted advice model (TAM) (Chattopadhyay *et al.*, 2016)—which provides actionable design guidelines for *trust-based* alerts. But how trust-based alerts affect physicians' clinical encounters remains unexplored.

Second, notwithstanding the critical importance of DDI alerts, understanding their effectiveness remains elusive. The effectiveness of clinical alerts is typically measured as the rate of 'canceling' the order instead of 'continuing'. Recent research suggests that such a binary outcome measure may be flawed because it fails to capture other positive effects of an alert, such as triggering preventive monitoring actions (Baysari *et al.*, 2016).

This paper evaluates trust-based alerts—by both measuring alert override rates and examining what other actions are triggered by alerts that may affect patient safety. Especially, we explore key questions about trust-based alerts, such as:

What trust cues, embedded in clinical alerts, can potentially promote reflection while prescribing medications? When and how do physicians respond to trust-based alerts in their decision-making? Moreover, how can we measure the effectiveness of alerts beyond traditional "cancel" or "continue" that considers long-term effects on patient conditions?

To that end, we conducted an explanatory sequential mixed-method study. Trust-based alerts improved adherence, promoted reflection, and triggered appropriate clinical actions. Findings indicate: (i) endorsed messages drive trust—but only if accompanied with data summaries, otherwise evoking condescension among physicians; (ii) situated awareness of adverse events (in the hospital or among prior patients) promotes reflection and triggers further monitoring, critical for patient safety; and (iii) future trust-based alerts should provide richer, alternative recommendations and link explanatory data resources beyond simple warnings and directives. Finally, we outline holistic metrics that go beyond the binary cancel/continue compliance measure and provide a broader perspective on how alerts can assist physicians in achieving patient safety.

## 2. RELATED WORK

DDI alerts provide clinical decision support (CDS) for health-care professionals prescribing medications to prevent adverse drug events (ADE) in patients (Jung *et al.*, 2012). DDI alerts are commonly interruptive, but other types are also used (Duke *et al.*, 2014). Upon receiving an interruptive DDI alert, physicians either cancel the drug order or continue by overriding the alert. Current override rates are alarmingly high, estimated as high as 95% (Bryant, Fletcher and Payne, 2014; van der Sijs *et al.*, 2009), which urges a fundamental rethinking of DDI alerts as a decision support system. DDI alerts are triggered from state-of-the-art knowledge bases (KBs). But as a practice, medicine lacks general rules that can be unambiguously applied to every case at hand (Hunter, 1996). Hence, although the medical field has shifted from anecdotal decision-making to evidence-based medicine (Sackett *et al.*, 1996), physicians often mistrust DDI alerts. This mistrust partly stems from a lack of alert specificity (Weingart *et al.*, 2009) and continues the influence of anecdotal evidence on medical practice. When making clinical judgments in atypical cases (Enkin and Jadad, 1998), physicians strongly believe in personal or trusted experiences over textbook rules (Greenhalgh, 1999; Jung *et al.*, 2012). But ignoring a clear majority of alerts (Bryant, Fletcher and Payne, 2014) greatly increases the risk of prescribing unsafe medications.

Research on safe prescribing practices indicates alert overload or alert fatigue as a primary cause of alert overrides (Bryant, Fletcher and Payne, 2014; Payne *et al.*, 2002). Physicians encounter too many alerts during their daily work to afford sufficient consideration. Prior work on improving physicians' consideration of clinical alerts explored alternative display strategies drawing on human-factor principles (Feldstein *et al.*, 2004; Russ *et al.*, 2014) and ways to advance knowledge bases to increase the positive predictive value of alerts (van der Sijs *et al.*, 2006). For instance, researchers examined alternative visual designs for alert content (Russ *et al.*, 2014), different positioning of alerts within the computerized provider order entry (CPOE) interface (Payne *et al.*, 2015; Wipfli *et al.*, 2016), contextual cues (Duke, Li and Dexter, 2013; Melton *et al.*, 2015), different temporal orders in the clinical workflow (Lo *et al.*, 2009; Shah *et al.*, 2006) and other human-factor issues (Russ *et al.*, 2014). For a non-interruptive experience, alerts are sometimes presented on a sidebar (Jung *et al.*, 2012) or regrouped at a location away from the prescription order entry form (Wipfli *et al.*, 2016). Overall, these solutions have had limited success in increasing alert compliance: regrouping alerts did not significantly change prescription behavior (Wipfli *et al.*, 2016), and alert adherence remained low after contextual cues were incorporated (~15%) (Duke, Li and Dexter, 2013). Besides presentational factors, researchers indicate that alerts are often ignored due to a lack of trust (Alexander, 2006; Hayward *et al.*, 2013; Zheng *et al.*, 2011). Achieving better specificity in alerts remains a challenge because of the little consensus on what alerts can be considered superfluous (Missiakos, Baysari and Day, 2015; van der Sijs *et al.*, 2006). Thus our goal is not to attain a specific percent compliance rate, but rather to ensure that the information being delivered to providers is received in a trustworthy manner for evaluation and clinical action. Although eliciting trust in instances of judgment uncertainty (Tversky and Kahneman, 1974) is a crucial requirement for DDI clinical decision support, current alerts have not been designed to embody trust cues. To that end, rather than improving alert KBs or presentational factors, we pursue a complementary approach and explore trust-eliciting cues drawing on aspects of trusted peer advice in face-to-face clinical settings (Chattopadhyay *et al.*, 2015).

### 3. METHODS

#### 3.1. Overview

The evaluative activities presented in this paper were situated in a larger design process (Fig. 1). The first phase of the project included a contextual inquiry to gather user requirements about trust-based design themes, which fed into design iterations. In consultation with domain experts, eleven final alert designs were created. Finally, those alert designs were evaluated in a two-part study, presented here.

#### 3.2. Designing DDI alerts

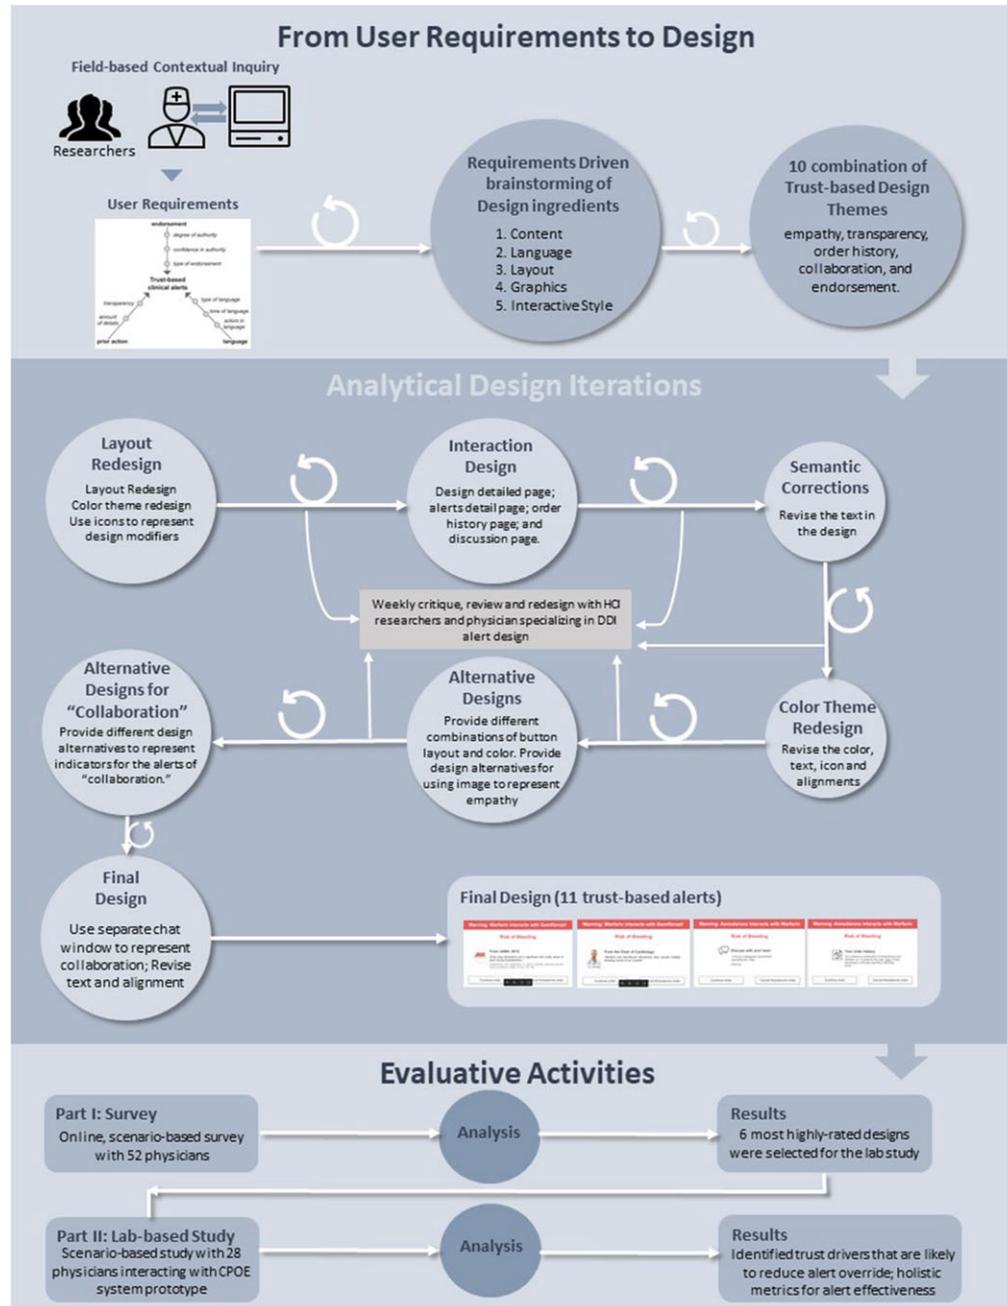
We design and evaluate DDI alerts incorporating different types of trust cues. These trust cues are informed by the TAM (Fig. 2) (Chattopadhyay *et al.*, 2016). TAM provides actionable design principles to incorporate trust into clinical alerts, based on three basic dimensions: different sources of *endorsement*, awareness of *prior actions*, and the use of a suitable language type (e.g. descriptive or prescriptive) and tone (e.g. negative or neutral). Each of these three dimensions has different sub-dimensions, which can be variously parameterized to design a wide range of trust-based clinical alerts. For instance, the sub-dimensions for endorsement are *degree of authority*, *confidence in authority* and *type of endorsement*. The function of TAM is to guide the design of trust-eliciting cues by parameterizing the different dimensions. Trust-eliciting cues for DDI alerts may be modeled using a combination of two, or three dimensions. However, any possible combination may not generate a reliable trust cue. For example, using authoritative language to portray overrides of another colleague would be considered directive or patronizing—not informative or helpful. Thus, first, design directions for trust cues need to be established, such as endorsed, transparent, or empathic alerts (Chattopadhyay *et al.*, 2015), and then the available dimensions parametrized to ideate design prototypes. The three trust-eliciting elements in TAM open up a rich design space for designing trust-eliciting cues in DDI alerts, thus not habituating physicians with alerts that look the same.

Grounded in prior work on trusted peer advice in clinical settings, we set out to establish strategies for embedding trust-eliciting cues in DDI alerts. Particularly, we explore two key questions:

- What trust cues, embedded in DDI alerts, increase alert effectiveness during prescribing medications?
- When and how do physicians respond to different types of trust cues in their clinical decision-making?

Ethnographic studies have examined how physicians structure their daily decisions on medication prescribing around mutual trust, through peer discussions and by attending to advice from their colleagues (Chattopadhyay *et al.*, 2015). The key elements of trusted peer advice are codified into the TAM (Chattopadhyay *et al.*, 2016) (Fig. 2). Using TAM, we designed a set of DDI alerts embodying a variety of trust cues and evaluated them in an explanatory sequential mixed-method study.

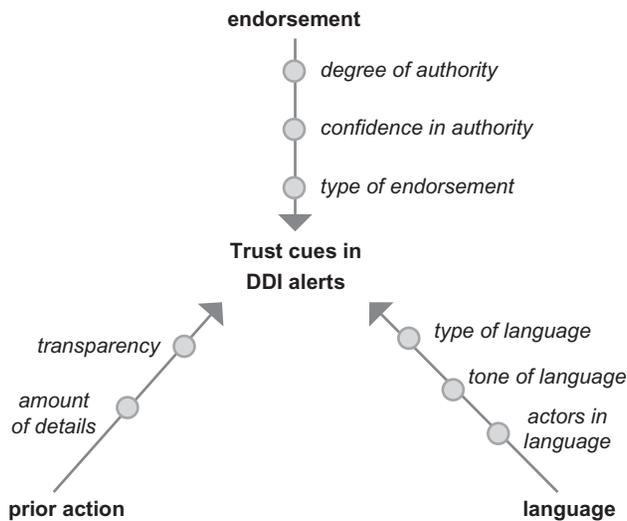
An accurate measure is a prerequisite for improvement. Research suggests a binary compliance measure is inadequate because it fails to assess all possible effects of encountering a DDI alert, such as monitoring actions, dosage modifications or provider-patient interactions (Hayward *et al.*, 2013; Russ *et al.*, 2012). Recent work recommended a Bayesian framework to identify physicians' override reasons into false



**Figure 1.** The evaluative activities presented in this paper were part of a larger design process, which included contextual inquiry to gather user requirements and iterative design activities in consultation with domain experts.

positives and false negatives to measure alert effectiveness (Payne *et al.*, 2015). Furthermore, studies have found physicians consult alerts only in rare, unfamiliar, and complex medical situations while ignoring them for routine prescriptions (Wipfli *et al.*, 2011). Although the override rate is widely used to measure alert effectiveness (McCoy *et al.*, 2012; Nanji *et al.*, 2014; Topaz *et al.*, 2015; van der Sijs *et al.*, 2006), it lacks *construct validity*. An alert may play an

effective role in clinical care, even if it was overridden; new tests monitoring patient conditions may be ordered, other medications modified or newly ordered, patient notes enriched, or expert clinical advice solicited. To operationalize and measure alert effectiveness as the alert override rate, thus, threatens the validity of the construct. Our mixed-method study measured the likelihood of alert compliance and alert usefulness in a survey and examined alert-prompted clinical



**Figure 2.** The TAM provides three fundamental dimensions to design trust cues in clinical alerts: endorsement, awareness of prior actions and a suitable language. These trust cues inform our redesign of DDI alerts.

actions relevant to patient safety in a scenario-based simulation.

DDI alerts were redesigned based on the three fundamental dimensions of the TAM: endorsement, prior action, and language (Chattopadhyay *et al.*, 2016). However, all possible permutations of different parameters and parameter values of TAM do not yield meaningful trust cues. Thus, a broad set of DDI alerts with trust-eliciting cues were generated and iteratively pruned with domain experts, finally converging into ten types of alerts (see Table 1). Examples of situations where this pruning occurred include representations of the endorsement that the clinical expert perceived too ‘pushy’ to positively sway the physician’s response; alert designs where the trust cues were overly prominent and risked to overshadow the message about the severity of the DDIs and adverse events. The DDI alert interface currently used in the Regenstrief Gopher (Duke *et al.*, 2014) served as the control. The 10 DDI and control alerts, rendered as JPEG figures, were used for the survey. For the simulation, six of the alerts were integrated into a research version of the Gopher system, t-EMR using HTML (Fig. 3). Those six alerts were dynamic, i.e. the hyperlinks showed pertinent facts and alert content changed based on patient scenarios. Facts used in the DDI alerts, like physicians’ drug-prescribing history or the expert physician did not reflect actual figures or real persons, respectively.

### 3.3. Evaluating redesigned DDI alerts

The redesigned DDI alerts were evaluated using an explanatory sequential mixed-method approach (Creswell and Clark, 2007). First, in a survey, 52 physicians reported the likelihood

of alert compliance and alert usefulness. Usefulness was operationalized as a five-item scale: convincing, trustworthy, annoying, helpful, and manipulative. Alerts were pruned and six alerts (including control) were selected based on the survey results for the scenario-based simulation, where 18 physicians used a CPOE to prescribe medications in six scenarios.

#### 3.3.1. Part I survey

The online survey constituted of seven questions. First, participants responded to 7-point Likert scales (from left to right: very, moderately, slightly, neutral, slightly, moderately, very) to six questions. Questions included alert compliance (‘How likely are you to continue this order?’) and usefulness (e.g. ‘This alert is *convincing*.’). Then, in an open-ended question, participants provided an explanation of their alert compliance response (e.g. ‘In your prior response, you indicated that you were very likely to override this alert. Could you briefly mention why?’). Pilot tests indicated a completion time of about 10 min.

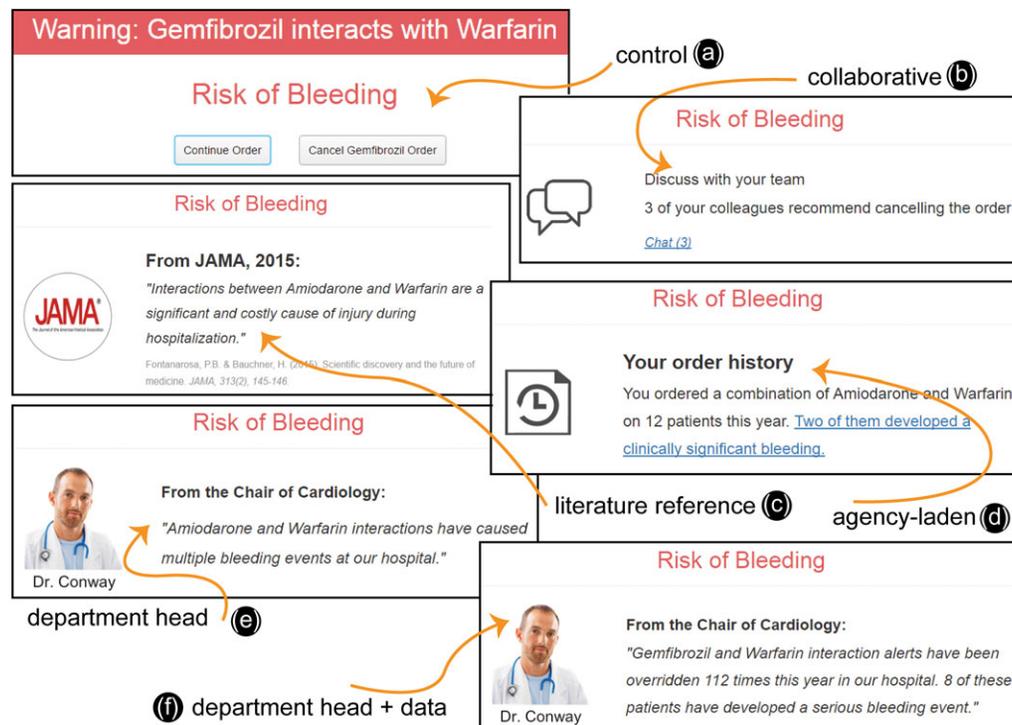
A call for participation with the survey link was sent out to physicians’ listservs in the Eskenazi Health Network during March 2016. Following informed consent and study instructions, participants were provided with two comparable medical scenarios. They were instructed to assume the role of a second-year medical resident in an inpatient General Medicine team and situate themselves in a typical medication prescribing setting to attend to the scenario at hand. The sequence of alerts in Group 1 was introduced by a medical scenario (see Supplementary File 1) and comprised of *control*, *endorsed*, *agency-laden*, *collaborative*, *empathy-driven*, and *transparent* alerts (Table 1). The type of endorsement for the endorsed alert was *department head + data*. Group 2 was also introduced by a medical scenario (see Supplementary File 1) and alerts comprised of six types of endorsed alerts, *department head + data*, *department head*, *department head + empathy*, *literature reference*, *Institute of medicine* and *chief of quality*. Alerts were completely randomized within the groups, but group order and scenario-to-group mapping remained identical across participants. Survey responses, response time and demographics were measured.

#### 3.3.2. Part II scenario-based simulation

The experimental setting simulated a room where physicians used a computer terminal to review patient charts and prescribe medications (Fig. 4). A tailored version of the Medical Gopher (Duke *et al.*, 2014) was used as the technology test bed. Each participant (recruited through the Eskenazi Health Network of physicians) encountered six patient scenarios with a different type of alert presented for each. The scenarios were designed by a clinician (JD) and reviewed with non-participating clinical colleagues for authenticity and consistency with typical inpatient clinical encounters. The order of alerts and scenarios were completely randomized across participants. When prescribing a medication as part of the task, a DDI alert (one of the six types) would trigger. The task was

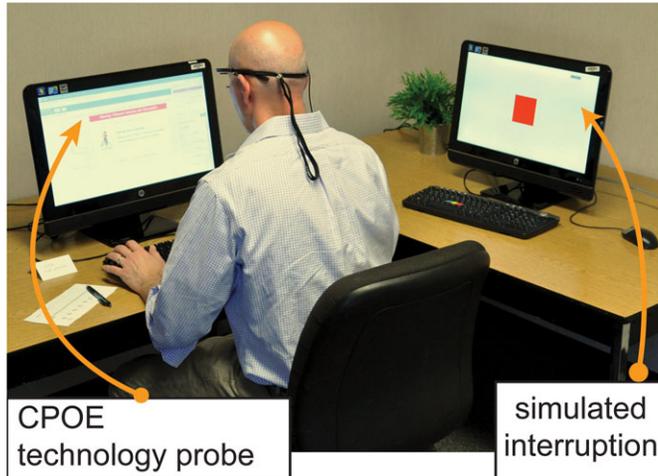
**Table 1.** Trust-eliciting cues generated using TAM to redesign DDI alerts.

	Type of alerts	Endorsement	Prior action	Language	
				Type	Actor
1.	Control	x	x	x	x
2.	Department head	Department head or specialist	x	Descriptive	Authority
3.	Department head + data	Department head or specialist	Data on patient population within the hospital	Descriptive	Authority
4.	Department head + empathy	Department head or specialist	x	Prescriptive	Authority
5.	Chief of quality	Chief of quality		Descriptive	Community
6.	Agency-Laden	x	Provider's prescribing history	Descriptive	Computer
7.	Collaborative	x	x	Prescriptive	Peers
8.	Empathy-driven	x	x	Descriptive	Computer
9.	Transparent	x	x	Descriptive, implying personal responsibility	Computer
10.	Literature reference	Literature	x	Descriptive	Peers
11.	Institute of medicine	Institute of medicine	x	Descriptive	Community

**Figure 3.** Six types of redesigned DDI alerts embodying different trust-eliciting cues: (a) control, (b) collaborative, (c) literature reference, (d) agency-laden, (e) department head and (f) department head + frequency of ADE in the hospital.

to review patient data, order a set of medications including continuing or canceling orders following the DDI alert and ordering any other appropriate tests. Participants used the CPOE for about 10–15 min.

To simulate interruptions, which are frequent in clinical settings, and add a realistic cognitive overload, a secondary experimental task was introduced (Wu *et al.*, 2014). On a second computer display (not within the participant's peripheral



**Figure 4.** DDI alerts were displayed in an experimental setting that simulated an inpatient meeting room where physicians use a workstation to prescribe medications. Interruptions were simulated using a second computer and screen interactions recorded using Camtasia.

vision), a gray square at the center of the screen changed to blue, yellow or red at random intervals (5, 10 or 15 s). The ‘g’, ‘h’ and ‘j’ keys of the QWERTY keyboard were replaced with blue, yellow and red colored keys. When the color of the square changed, participants had to respond by pressing the matching colored key, as quickly and as accurately as possible. They were asked to monitor for any change of color and deem both tasks as equally important.

After the simulation, alert screenshots were used to conduct a stimulated recall interview. Participants recalled their behavior toward each medical scenario and explained their decision-making rationale. The entire study took about an hour and 15 min. Screen interactions were captured using Camtasia and interviews audio recorded. Before the study, a facilitator introduced the CPOE system and explained the color-matching task. Patient scenarios and alert designs used in the survey and simulation are available in Supplementary Files 1 and 2, respectively. Each participant was given \$85 for their participation. This study was approved by IU IRB no. 1 509 280 592.

## 4. RESULTS

### 4.1. Survey

#### 4.1.1. Quantitative analysis

Data did not follow parametric assumptions; the Shapiro-Wilk test was significant,  $P < 0.05$ . Data were analyzed with Friedman’s ANOVA and *post hoc* pairwise comparisons with Wilcoxon signed-rank test and Bonferroni correction. Some responses to the alert compliance question were lost due to a database failure.

#### 4.1.2. Findings

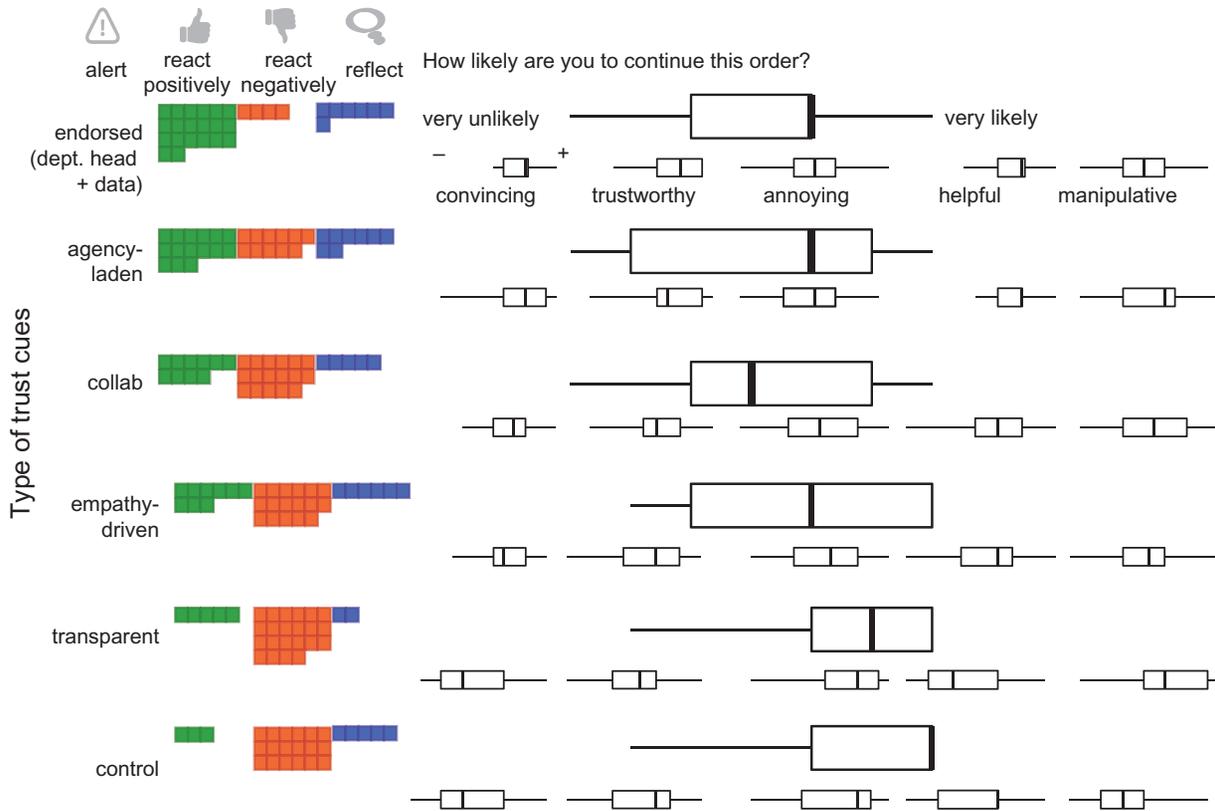
52 physicians (20 females,  $Mdn_{age} = 28$ ,  $IQR = 2.25$ ) from Eskenazi Health participated in this study; 2 were attending physicians, 15 interns, 2 others and 33 residents. Most were experienced with CPOE, with 32 participants using CPOE for 1–3 years and 13 for 4 or more years; 48 spent 50% or more of their time in an inpatient environment.

In Group 1, the alert type significantly affected physicians’ likelihood of compliance,  $\chi^2(5) = 45.91$ ,  $P < 0.001$ ,  $n = 29$ , and perceived usefulness,  $n = 52$ ,  $P < 0.001$  (Fig. 5). *Post hoc* tests indicated physicians were significantly less likely to override agency-laden, endorsed, and collaborative alert than control,  $P < 0.001$ . Endorsed alert was significantly more convincing than control and transparent,  $P < 0.001$ , more trustworthy than transparent, empathy-driven and control,  $P < 0.001$ , more helpful than transparent and control,  $P < 0.001$ , but less annoying than transparent alert,  $P < 0.001$ . Agency-laden alert was significantly more convincing than control and transparent,  $P < 0.001$ , more trustworthy than control,  $P < 0.001$ , more helpful than control and transparent,  $P < 0.001$ , and more manipulative than control,  $P < 0.001$ ; but less annoying than the transparent alert,  $P < 0.001$ . The collaborative alert was significantly more convincing than control and transparent,  $P < 0.001$ , more helpful than transparent,  $P < 0.001$ , and more manipulative than control,  $P < 0.001$ . Differences in response times were not significant.

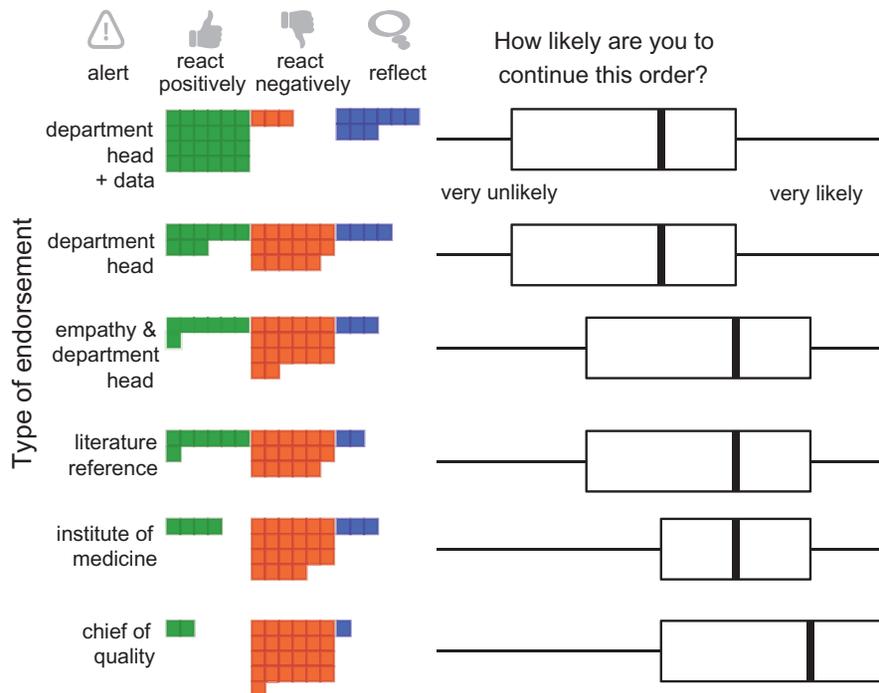
Similarly, in Group 2, the type of endorsement significantly affected physicians’ likelihood of compliance,  $\chi^2(5) = 34.49$ ,  $P < 0.001$ ,  $n = 29$ , and usefulness,  $n = 52$ ,  $P < 0.001$  (Fig. 6). *Post hoc* tests indicated physicians were significantly less likely to override endorsement from department head + data than the chief of quality,  $P < 0.001$ . Department head + data was significantly more helpful than all alerts,  $P < 0.001$ , more convincing than all alerts, except department head,  $P < 0.001$ , and more trustworthy than the chief of quality and institute of medicine,  $P < 0.001$ , but less annoying than the chief of quality,  $P < 0.001$ .

#### 4.1.3. Qualitative analysis

Qualitative responses explaining physicians’ likelihood of compliance were coded ( $n = 29$ ) into three overarching themes: *positive* reaction, *negative* reaction and *reflection*. A response could be positive, negative or neither (mutually exclusive), and optionally, demonstrate reflection (Figs 4 and 5). Positive reaction suggested alerts were perceived more trustworthy, convincing and helpful, but less annoying and manipulative than the control. Further analysis indicated that positive reactions were prompted by expert endorsements with data; on the other hand, the lack of data and sweeping remarks produced negative feelings of condescension, futility and manipulation. Descriptive statistics about the hospital and cases of prior patients heightened the value of endorsements.



**Figure 5.** Physicians were significantly more likely to override the *control* alert than *agency-laden*, *endorsed* and *collaborative* alert,  $P < 0.001$ . These three types of alerts were selected for the scenario-based simulation study.



**Figure 6.** Type of endorsement significantly affected physicians' intended compliance,  $P < 0.001$ . Endorsement from the *department head with data* was significantly more helpful than all other alerts,  $P < 0.001$ .

#### 4.1.4. Findings

Expert endorsement, computerized advice from ‘known and reliable sources’, were necessary, but not sufficient in eliciting trust. For instance, generic remarks from the chief of quality or institute of medicine prompted disdain (P3: ‘*Who is this guy? Generic statement no clinical context*’; P14: ‘*I am distracted by the canned statement and forgot about the bleeding risk*’), condescension (P11: ‘*feels patronizing and it’s extra text that we don’t gain anything from reading*’) or manipulation (P5: ‘*Feels like playing to our emotions [...]*’). Overall, blanket statements were found annoying and doubting physicians’ commitment to patient safety (P15: ‘*I’m not sure why the ‘chair’s’ opinion would change my clinical judgment unless I asked for a consult.*’). However, DDI alerts accompanied with situated data, such as hospital statistics of ADE or physicians’ prior actions promoted reflection (P14: ‘*Seeing the outcomes of my peers and my clinical care makes it more likely that I would pause and consider whether this medication combination was the correct one.*’ P27: ‘*Comes from a known person and has stats pertinent to my workplace.*’). Department head + data evoked most positive reactions compared with other endorsed alerts. Notably, literature reference evoked mixed feelings (P5: ‘*Would like specific numbers from JAMA before I change my decision but I like reference indicated and reliable source.*’ P9: ‘*This is more trustworthy and may make me think more but I doubt I would have time to read the trial.*’)

Transparent alerts provoked resistance toward being surveilled (P27: ‘*I don’t respond well to threats. Sue me.*’ P5: ‘*This feels like the [alert] is trying to place blame on a provider, [...]. [It] is set up for a future malpractice suit.*’). Even collaborative alerts were sometimes perceived as annoying (P15: ‘*Having to explain the clinical rationale to colleagues is a bit too much ‘big brother.’ Perhaps this clinical decision, is in fact, the correct one but answering to peers is frustrating as if you are doing something wrong when in fact the dosage for these medications can be adjusted to make this a safe combination.*’)

The survey results informed the selection of five trust-based alerts for further exploration in the simulation: *collaborative, agency-laden, literature reference, department head, and department head + data.*

## 4.2. Simulation

### 4.2.1. Participants

About 18 physicians (10 females,  $Mdn_{age} = 27.5$ ,  $IQR = 1$ ) from Eskenazi Health participated in this study; 2 were interns, 15 residents and 1 attending provider. Participants were moderately familiar with a CPOE system ( $Mdn = 6.5/10$ ,  $IQR = 1$ ) and responded to the simulated interruption with an average efficiency of 3.62 s ( $SD = 7.72$ ) and 84.35% accuracy ( $SD = 18.51$ ).

### 4.2.2. Procedure

Following the simulation, participants ( $n = 18$ ) recalled their experience with DDI alerts (Fig. 3) and explained their decision rationale. Using a laddering technique (Reynolds and Gutman, 1988), participants discussed when, why and what information in the alerts was helpful, trustworthy and convincing. Random combinations of the alert and scenario generated a variety of DDI encounters—providing the opportunity to qualitatively assess alert-inducing behavior in a range of different situations.

### 4.2.3. Qualitative analysis

Qualitative analysis was conducted collaboratively by two researchers through inductive, open coding of the themes that emerged in the stimulated recall interviews. The themes demonstrated physicians’ experiential thought processes while interacting with the trust-based alerts during decision-making. Themes and sub-themes from the participants’ experience and reactions were also visually mapped back to the six alerts designs. This experience-to-design analysis allowed researchers to gain a clearer understanding of the design factors that participants were referring to in their interview responses. The findings elucidated three high-level themes that characterize the physicians’ experience with trust-based alerts.

### 4.2.4. Been there, done that: overriding by experience

Alerts were mostly overridden—often without further reflection—when physicians encountered familiar situations. For example, several of our participants expressed that the warning about Warfarin was extremely common, which reduces the potency of the alert in causing any reflection:

‘It really doesn’t scare me that much, because we get alerts like this constantly. Everything interacts with warfarin.’ (P18)

‘Honestly, just in the hospital, from seeing so many patients on gemfibrozil and warfarin, and warfarin basically interacting with everything, so just because we do it all the time, I ignored the interaction.’ (P13)

Similarly, another reason behind noncompliance had to do with the gap between their experiential knowledge and the alerts. Participants repeatedly mentioned their decision rationale to override alerts as ‘not seen in practice’. One participant here demonstrates the difference between what the EMR systems alert him about from his experience and how that feeds into his decision when prescribing:

‘...they always say risk of rhabdo, risk of rhabdo, and you never see it, and not just like with diltiazem, also with others like [...] So just because I have never seen it.’ (P13)

The skepticism towards alert warnings our participants experienced was also extended to endorsed alerts by department head when no data summaries were presented. For instance, for P18, just an endorsement from an authority figure didn’t trump his own experience:

'I feel like I have seen diltiazem withstands [holds up] all the time, at least frequently. [.]' (P18)

However, when facing unfamiliar situations, DDI alerts altered provider behavior, triggering further monitoring of patient vitals, and promoting reflection. Here the same participant reflects on an instance of prescribing a drug he does not have experience with:

'It [alert] did [influence] my decision. That's partly because I don't use dihydro or itraconazole very often. So, it was an unfamiliar medication regard to the side effects. It made me wanna look up the side effects of azoles to know what the toxicity is.' (P18)

Here the physician found the alert to be helpful while faced with uncertainty in his clinical decision-making. By providing information for further investigation, the trust-eliciting cues facilitated evaluating risk-benefit concerns.

#### 4.2.5. Alerts prompt weighing of risk over benefits

Based on our findings, when evaluating risk-benefit concerns, often physicians opted for rigorous monitoring, rather than holding off on prescribing a certain medication. For example, when alerts did induce further reflection, several participants explained that the alerts serve the purpose of creating awareness and/or further testing, but not halting the medication:

'It made me wanna get more labs rather than not order the medication.' (P18)

'Just would monitor for the rhabdo.' (P17)

'You want to be aware of it if something does happen, but I didn't think it [medication] was necessary to hold off.' (P16)

Thus, rather than canceling the order at hand, an alert's impact manifested in additional monitoring or deciding on which of the two reacting drugs to stop or modify by '*kind of weighing the risk-benefit of (the drugs)*' (P13). Changes to prescription orders were prompted when some ADE were more severe than others. Here, our participants demonstrate how the alerts prompted the negotiation process in weighing risks and benefits for their patients:

'Yes, it made me change what I was going to prescribe. Just because QT prolongation scares me.' (P17)

'Headaches are uncomfortable, but not life-threatening, but you may induce a bad drug toxicity because you have not explored other options.' (P3)

Trust-eliciting cues in the alerts created opportunities for physicians to consider the alert with specific patient and risk-driven factors.

#### 4.2.6. The role of situated data in increasing awareness

Consistent with the survey results, DDI alerts elicited most positive reactions when data summaries—personal to the provider—were displayed. For example, key facts on ADE that

occurred in the physicians' hospital or among their prior patients, played a key role in promoting reflection and further monitoring or changes in the order. Here a couple of participants reflect on the cues that supported their decision to comply with the alert:

'The most helpful ones were the ones that stated what had happened in this hospital or what have been your personal experience.' (P16)

'This usually does bother me whenever I have had a patient that has had a bad reaction to medications that I have prescribed. It's definitely different whenever it is the computer telling me that it happens than in real life.' (P18)

Because those alerts provided situated data, they prompted additional awareness for physicians. Physicians indicated feeling a push to consider and find alternative prescriptions. Similarly, this kind of personalized alerts also caused alarm in our participants:

'This one's the most alarming alert because like it's from somebody from the hospital means it has happened in the hospital. It'd have been scarier if you get alert like 'hey last month 10 of the patients got like severe adverse reaction because of what you just did'.' (P13)

Through these cues, physicians indicated increased awareness about organizational efforts in improving patient safety

'From the chair of cardiology. So, it is something the hospital is working on to improve quality. So, we could find a different antibiotic that is more suitable.' (P16)

Alerts embodying local, organizational data induced further consideration as physicians became more aware of how their actions feed into organization-wide metrics.

## 5. DISCUSSION

Findings indicated that (i) endorsed messages drive trust when accompanied with data summaries, otherwise evoke condescension, and (ii) situated awareness about the frequency of ADE (in the hospital or among prior patients) promotes reflection and triggers further monitoring, which is critical for patient safety. Table 2 lists our design recommendations for DDI alerts in the context of prior work in this area.

### 5.1. Recommendations for designing DDI alerts to elicit trust

Guidelines for standardized reporting of DDI recommend describing the incidence of each ADE and the clinical impact of the DDI at the population level, such as frequency and severity of incidences (Floor-Schreuderling *et al.*, 2014). Design recommendations for DDI alerts suggest displaying the ADE frequency, when available (Payne *et al.*, 2015).

**Table 2.** Design recommendations to improve DDI alert effectiveness.

Design requirement	Design recommendation	Related guideline(s)
Combat illusory correlations <sup>a</sup>	<ul style="list-style-type: none"> <li>• Provide frequency of ADE in the hospital and link descriptions of incidences</li> <li>• Provide frequency of ADE in prior patients of the interacting provider and link descriptions of incidences, source: findings from the simulation (Results: Simulation, Been There, Done That)</li> </ul>	Prior guidelines recommended displaying global ADE frequency in alerts, when available (Payne <i>et al.</i> , 2015) and reporting ADE for DDI management (Floor-Schreuderling <i>et al.</i> , 2014).
Facilitate accuracy goal <sup>b</sup>	<ul style="list-style-type: none"> <li>• Provide endorsement from local medical specialists, such as the chair of cardiology, or chief pharmacist, source: findings from survey and simulation (Results: Simulation, Been There, Done That &amp; Survey, Quantitative Findings)</li> </ul>	
Enrich personal knowledge base	<ul style="list-style-type: none"> <li>• Allow physicians to bookmark alerts for later review of all the displayed information, in a less time-constrained situation, source: findings from the simulation (Results: Simulation, Been There, Done That).</li> </ul>	Provide focused educational materials during medication prescribing in the CPOE, e.g. a summary of disease-specific national guidelines or links to educational monographs, displayed on a side pane (Miller <i>et al.</i> , 2005).
Tailor displayed information based on clinical role and frequent/infrequent alerts	<ul style="list-style-type: none"> <li>• Modify amount of information displayed on the alert based on physicians' job experience and clinical role. More experienced physicians should encounter less information, while details remain available on demand (e.g. embedded links or inline summaries), source: findings from the simulation (Results: Survey, Qualitative Findings).</li> </ul>	Alert response times among providers with different clinical roles were significantly different (McDaniel <i>et al.</i> , 2016) as was for frequent vs. infrequent alerts in a hospital; tailoring of alerts is suggested for individual physicians and different medical specialties (Kesselheim <i>et al.</i> , 2011).

<sup>a</sup>Illusory correlations occur when physicians make memory-based judgments, evaluating the probability of an ADE by the ease with which similar instances can be recounted (Chaiken *et al.*, 1989; Chapman & Chapman, 1969).

<sup>b</sup>Making judgments under time pressure lead to hasty reasoning, exaggerating cognitive biases and heuristic processing using simple rules. When motivated to be accurate, people spend more cognitive effort in reasoning, consider pertinent information carefully, and process it more deeply and systematically, using complex rules (Kunda, 1990).

Confirming earlier studies, we provide empirical evidence, based on the simulation study, that information on ADE frequency elicit trust. Our results further establish a role of expert endorsements and situated awareness in driving trust in DDI alerts. However, without any facts, advice from experts or authority figures presenting clinical judgment, showing empathy, or reminding institutional objectives was perceived as canned statements. A simple reminder that alert overrides would be recorded and viewable by other clinicians (Payne *et al.*, 2015) (transparent alert) was considered too 'big brother' and failed to positively affect alert behavior.

These findings show the limitations of anthropomorphizing computers to elicit emotions, such as trust (Culley and Madhavan, 2013; Hall and Cooper, 1991; Perse *et al.*, 1992). When conveyed through a computer, what may seem like an innocuous judgment call in person (expert opinion) was considered irksome unsolicited advice; similarly, a reminder was considered a threat inducing fear, guilt, liability and feelings of unfair surveillance (transparent alert). In sum, the trust factors uncovered from face-to-face clinical settings (Chattopadhyay *et al.*, 2015) that were successful in clinical

decision support (CDS) were—impersonal data and personal contexts.

In implementing these findings, we do not recommend that every DDI alert is endorsed. Rather, department administrators may decide to add an endorsement to certain alerts, such as high-priority DDIs (McEvoy *et al.*, 2016)—to promote providers' consideration. Depending on quality control outcomes or population health data in a hospital, either hospital statistics on ADE or individual physicians' practice statistics, or both, may be displayed in an alert. While specifying the frequency of ADE, when available, have been suggested previously (Payne *et al.*, 2015), we argue that rather than global factors (e.g. literature reference), situated awareness is crucial in driving physicians' trust.

Selective customization of DDI alerts could further mitigate the concerns of alert override as a habitual behavior (Baysari *et al.*, 2016). Tailoring warnings for a particular clinical environment, relevant to a patient's demographics or medication dosage, or different specialties was previously recommended to minimize warnings and thereby reduce alert fatigue (Kesselheim *et al.*, 2011). Although current Electronic Health

Record (EHR) vendors severely limit a hospital's ability to customize alert systems in fear of liability, a recent review of product liability principles indicated tailored alerts do not raise manufacturers' and physicians' risk of liability, but could help lower both the liability risk and perceptions of risk (Kesselheim *et al.*, 2011). It was evident from the simulation that tailoring DDI alerts to facilitate situated awareness would make them more trustworthy and thus, effective in upholding safe prescribing practices.

## 5.2. Other design recommendations

Other than trust-eliciting cues, several DDI alert design requirements emerged during the study. Some of those echoed prior design recommendations, such as embedding links to pertinent literature (e.g. research highlights of a JAMA article) or web resources (e.g. UpToDate), DDI severity tiering, providing actionable alternative strategies, and offering information on the mechanism of the DDI (McEvoy *et al.*, 2016; Payne *et al.*, 2015). We also found from the simulation part of the study that physicians were predisposed to use the information presented in a DDI alert as elements to structure their personal knowledge base, which may reflect the demographic majority of our study—residents in their late twenties. Our results revealed that physicians almost always valued alerts and browsed associated information when the alerts were unfamiliar to them (or rarely encountered; e.g. 'I didn't know about that interaction.'). Because of a wide variety of expertise and experience, we recommend supplementary information on DDI alerts be tailored—according to specialty and/or clinical role (i.e. attending vs. intern). Recommendation for tailoring DDI alerts is not entirely new. Prior work showed that alert response times significantly vary among providers with different clinical roles (McDaniel *et al.*, 2016), as well as for frequent vs. infrequent alerts in a hospital. Previously, researchers suggested tailoring alerts for individual physicians and different medical specialties (Kesselheim *et al.*, 2011). While these works deal with whether to trigger or suppress an alert for a certain provider, we recommend tailoring *the supplemental information in a DDI alert* according to clinical role and frequency of occurrence in a hospital. The potential for in-situ collaborative opportunities with specialists, such as pharmacists, was considered beneficial based on the survey and simulation responses, but overall, responses were mixed.

During post-test interviews, often participants remarked that more information on the alert would be helpful for decision-making. But they rarely attended to the additional information already available on the alert interface (displayed or linked) while completing an order. We suggest tailoring information presented in an alert based on a provider's title (e.g. intern, resident or an attending) and specialty. Because experience plays a significant role in medicine, an intern would need more resources to make an informed decision

than an attending, but statistics on recent hospital events would be equally helpful to both. Table 2 lists our recommended design guidelines.

## 5.3. Limitations

This study has several important limitations. The implementation of the trust cues did not include real data. For example, the actual frequency of ADE in the hospital where study participants practiced, or the descriptions of different experts were fictitious. Patient scenarios were crafted based on realistic patient data from a teaching EMR system (t-emr) to make sure they could trigger relevant DDI alerts. In terms of deployment in hospital EMRs, we also acknowledge the challenge of attributing the cause of adverse events data to specific drug combinations as well as the intensive nature of updating the ever-evolving EHR knowledge bases.

## 6. CONCLUSIONS

We identified several trust-eliciting cues for improving the efficacy of DDI alerts. Trust factors did not directly translate from human-human interactions (Chattopadhyay *et al.*, 2015) to human-computer interactions in clinical settings. Expert endorsement of DDI alerts and situated awareness of ADE are recommended during the alert presentation. Situated awareness can prevent illusory correlations (Chaiken *et al.*, 1989) and nudge systematic processing during decision-making. DDI alerts embodying trust-eliciting cues can improve alert effectiveness, by increasing alert compliance or prompting appropriate clinical actions. These implications are important for interaction design researchers and practitioners to inform future research on the appropriate design and evaluation approaches for such alerts. Our recommendations also provide healthcare organizations and EHR vendors with important directions to improve alert design beyond presentational elements, to promote trust in CDS and thereby help physicians make an informed decision around patient safety.

## SUPPLEMENTARY MATERIAL

Supplementary data are available at *Interacting with Computers* online.

## ACKNOWLEDGEMENTS

We thank all our provider participants for their time, Yuan Jia for assisting with the visual design, Pankaj Avhad (IUPUI), Josh Castagno, Jeremy Leventhal and Haritha Mannam (Regenstrief Institute, Inc., Indianapolis, IN, USA) for developing and maintaining the experimental applications and website used in this study.

## FUNDING

This research is based upon work supported by the National Science Foundation Grant IIS-1343973. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Science Foundation.

## REFERENCES

- Alexander, G.L. (2006) Issues of trust and ethics in computerized clinical decision support systems. *Nurs. Adm. Q.*, 30, 21–29.
- Baysari, M.T., Tariq, A., Day, R.O. and Westbrook, J.I. (2016) Alert override as a habitual behavior—a new perspective on a persistent problem. *J. Am. Med. Inform. Assoc.*, ocv072.
- Bryant, A.D., Fletcher, G.S. and Payne, T.H. (2014) Drug interaction alert override rates in the meaningful use era: no evidence of progress. *Appl. Clin. Inform.*, 5, 802–813.
- Chaiken, S., Liberman, A. and Eagly, A.H. (1989). Heuristic and systematic information processing within and. *Unintended thought*, 212, 212–252.
- Chapman, L.J. and Chapman, J.P. (1969) Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *J. Abnorm. Psychol.*, 74, 271.
- Chattopadhyay, D., Duke, J.D. and Bolchini, D. (2016). Endorsement, Prior Action, and Language: Modeling Trusted Advice in Computerized Clinical Alerts. In Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (pp. 2027–2033). ACM.
- Chattopadhyay, D., Ghahari, R.R., Duke, J. and Bolchini, D. (2015) Understanding advice sharing among physicians: towards trust-based clinical alerts. *Interact. Comput.*, iww030.
- Creswell, J.W. and Clark, V.L.P. (2007). Designing and conducting mixed methods research.
- Culley, K.E. and Madhavan, P. (2013) A note of caution regarding anthropomorphism in HCI agents. *Comput. Human. Behav.*, 29, 577–579.
- Duke, J.D., Li, X. and Dexter, P. (2013) Adherence to drug–drug interaction alerts in high-risk patients: a trial of context-enhanced alerting. *J. Am. Med. Inform. Assoc.*, 20, 494–498.
- Duke, J.D., Morea, J., Mamlin, B., Martin, D.K., Simonaitis, L., Takesue, B.Y. and Dexter, P.R. (2014) Regenstrief Institute’s medical gopher: a next-generation homegrown electronic medical record system. *Int. J. Med. Inform.*, 83, 170–179.
- Enkin, M.W. and Jadad, A.R. (1998) Using anecdotal information in evidence-based health care: heresy or necessity? *Ann. Oncol.*, 9, 963–966.
- Feldstein, A., Simon, S.R., Schneider, J., Krall, M., Laferrriere, D., Smith, D.H. and Soumerai, S.B. (2004) How to design computerized alerts to ensure safe prescribing practices. *Jt. Comm. J. Qual. Saf.*, 30, 602–613.
- Floor-Schreudering, A., Geerts, A.F., Aronson, J.K., Bouvy, M.L., Ferner, R.E. and De Smet, P.A. (2014) Checklist for standardized reporting of drug–drug interaction management guidelines. *Eur. J. Clin. Pharmacol.*, 70, 313–318.
- Greenhalgh, T. (1999) Narrative based medicine in an evidence based world. *BMJ.*, 318, 323.
- Hall, J. and Cooper, J. (1991) Gender, experience and attributions to the computer. *J. Educ. Comput. Res.*, 7, 51–60.
- Hayward, J., Thomson, F., Milne, H., Buckingham, S., Sheikh, A., Fernando, B. and Pinnock, H. (2013) ‘Too much, too late’: mixed methods multi-channel video recording study of computerized decision support systems and GP prescribing. *J. Am. Med. Inform. Assoc.*, 20, e76–e84.
- Hunter, K. (1996) ‘Don’t think zebras’: uncertainty, interpretation, and the place of paradox in clinical education. *Theor. Med. Bioeth.*, 17, 225–241.
- Isaac, T., Weissman, J.S., Davis, R.B., Massagli, M., Cyrulik, A., Sands, D.Z. and Weingart, S.N. (2009) Overrides of medication alerts in ambulatory care. *Arch. Intern. Med.*, 169, 305–311.
- Jung, M., Riedmann, D., Hackl, W.O., Hoerbst, A., Jaspers, M.W., Ferret, L. and Ammenwerth, E. (2012) Physicians’ perceptions on the usefulness of contextual information for prioritizing and presenting alerts in computerized physician order entry systems. *BMC. Med. Inform. Decis. Mak.*, 12, 111.
- Kesselheim, A.S., Cresswell, K., Phansalkar, S., Bates, D.W. and Sheikh, A. (2011) Clinical decision support systems could be modified to reduce ‘alert fatigue’ while still minimizing the risk of litigation. *Health. Aff.*, 30, 2310–2317.
- Kunda, Z. (1990) The case for motivated reasoning, 108. *Psychol. Bull.*, 480, 480.
- Lo, H.G., Matheny, M.E., Seger, D.L., Bates, D.W. and Gandhi, T. K. (2009) Impact of non-interruptive medication laboratory monitoring alerts in ambulatory care. *J. Am. Med. Inform. Assoc.*, 16, 66–71.
- McCoy, A.B., Waitman, L.R., Lewis, J.B., Wright, J.A., Choma, D. P., Miller, R.A. and Peterson, J.F. (2012) A framework for evaluating the appropriateness of clinical decision support alerts and responses. *J. Am. Med. Inform. Assoc.*, 19, 346–352.
- McDaniel, R.B., Burlison, J.D., Baker, D.K., Hasan, M., Robertson, J., Hartford, C. and Hoffman, J.M. (2016) Alert dwell time: introduction of a measure to evaluate interruptive clinical decision support alerts. *J. Am. Med. Inform. Assoc.*, 23, e138–e141.
- McEvoy, D.S., Sittig, D.F., Hickman, T.T., Aaron, S., Ai, A., Amato, M. and Krall, M.A. (2016) Variation in high-priority drug–drug interaction alerts across institutions and electronic health records. *J. Am. Med. Inform. Assoc.*, 24, 331–338.
- Melton, B.L., Zillich, A.J., Russell, S.A., Weiner, M., McManus, M. S., Spina, J.R. and Russ, A.L. (2015) Reducing prescribing errors through creatinine clearance alert redesign. *Am. J. Med.*, 128, 1117–1125.
- Miller, R.A., Waitman, L.R., Chen, S. and Rosenbloom, S.T. (2005) The anatomy of decision support during inpatient care provider order entry (CPOE): empirical observations from a decade of CPOE experience at Vanderbilt. *J. Biomed. Inform.*, 38, 469–485.
- Missiakos, O., Baysari, M.T. and Day, R.O. (2015) Identifying effective computerized strategies to prevent drug–drug interactions in hospital: a user-centered approach. *Int. J. Med. Inform.*, 84, 595–600.

- Nanji, K.C., Slight, S.P., Seger, D.L., Cho, I., Fiskio, J.M., Redden, L.M. and Bates, D.W. (2014) Overrides of medication-related clinical decision support alerts in outpatients. *J. Am. Med. Inform. Assoc.*, 21, 487–491.
- Payne, T.H., Hines, L.E., Chan, R.C., Hartman, S., Kapusnik-Uner, J., Russ, A.L. and Glassman, P.A. (2015) Recommendations to improve the usability of drug-drug interaction clinical decision support alerts. *J. Am. Med. Inform. Assoc.*, 22, 1243–1250.
- Payne, T.H., Nichol, W.P., Hoey, P. and Savarino, J. (2002). Characteristics and Override Rates of Order Checks in a Practitioner Order Entry System. In *Proceedings of the AMIA Symposium* (p. 602). American Medical Informatics Association.
- Perse, E.M., Burton, P.I., Lears, M.E., Kovner, E.S. and Sen, R.J. (1992) Predicting computer-mediated communication in a college class. *Commun. Res. Rep.*, 9, 161–170.
- Reynolds, T.J. and Gutman, J. (1988) Laddering theory, method, analysis, and interpretation. *J. Advert. Res.*, 28, 11–31.
- Russ, A.L., Zillich, A.J., McManus, M.S., Doebbeling, B.N. and Saleem, J.J. (2012) Prescribers' interactions with medication alerts at the point of prescribing: a multi-method, in situ investigation of the human-computer interaction. *Int. J. Med. Inform.*, 81, 232–243.
- Russ, A.L., Zillich, A.J., Melton, B.L., Russell, S.A., Chen, S., Spina, J.R. and Hawsey, J.M. (2014) Applying human factors principles to alert design increases efficiency and reduces prescribing errors in a scenario-based simulation. *J. Am. Med. Inform. Assoc.*, 21, e287–e296.
- Sackett, D.L., Rosenberg, W.M., Gray, J.M., Haynes, R.B. and Richardson, W.S. (1996). Evidence based medicine: what it is and what it isn't.
- Shah, N.R., Seger, A.C., Seger, D.L., Fiskio, J.M., Kuperman, G.J., Blumenfeld, B. and Gandhi, T.K. (2006) Improving acceptance of computerized prescribing alerts in ambulatory care. *J. Am. Med. Inform. Assoc.*, 13, 5–11.
- Topaz, M., Seger, D.L., Lai, K.H., Wickner, P.G., Goss, F.R., Dhopeswarkar, N. and Zhou, L. (2015) High override rate for opioid drug-allergy interaction alerts: current trends and recommendations for future. *Medinfo.*, 216, 242–246.
- Tversky, A. and Kahneman, D. (1974) Judgement under uncertainty: heuristics and biases. *Sci. New. Ser.*, 185, 1124–1131.
- van der Sijs, H., Aarts, J., Vulto, A. and Berg, M. (2006) Overriding of drug safety alerts in computerized physician order entry. *J. Am. Med. Inform. Assoc.*, 13, 138–147.
- van der Sijs, H., Aarts, J., van Gelder, T., Berg, M. and Vulto, A. (2008) Turning off frequently overridden drug alerts: limited opportunities for doing it safely. *J. Am. Med. Inform. Assoc.*, 15, 439–448.
- van der Sijs, H., Mulder, A., van Gelder, T., Aarts, J., Berg, M. and Vulto, A. (2009) Drug safety alert generation and overriding in a large Dutch university medical centre. *Pharmacoepidemiol. Drug. Saf.*, 18, 941–947.
- Weingart, S.N., Massagli, M., Cyrulik, A., Isaac, T., Morway, L., Sands, D.Z. and Weissman, J.S. (2009) Assessing the value of electronic prescribing in ambulatory care: a focus group study. *Int. J. Med. Inform.*, 78, 571–578.
- Wipfli, R., Betrancourt, M., Guardia, A. and Lovis, C. (2011) A qualitative analysis of prescription activity and alert usage in a computerized physician order entry system. *Stud. Health. Technol. Inform.*, 169, 940–944.
- Wipfli, R., Ehrler, F., Bediang, G., Bétrancourt, M. and Lovis, C. (2016) How regrouping alerts in computerized physician order entry layout influences physicians' prescription behavior: results of a crossover randomized trial. *JMIR Hum. Factors*, 3. doi: 10.2196/humanfactors.5320.
- Wu, L., Cirimele, J., Leach, K., Card, S., Chu, L., Harrison, T.K. and Klemmer, S.R. (2014). Supporting Crisis Response with Dynamic Procedure Aids. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (pp. 315–324). ACM.
- Zheng, K., Fear, K., Chaffee, B.W., Zimmerman, C.R., Karls, E.M., Gatwood, J.D. and Pearlman, M.D. (2011) Development and validation of a survey instrument for assessing prescribers' perception of computerized drug-drug interaction alerts. *J. Am. Med. Inform. Assoc.*, 18, i51–i61.